Administration Guide

hp StorageWorks NAS 4000s and 9000s

First Edition (November 2003)

Part Number: 352405-001

This guide provides information on performing the administrative tasks necessary to manage the HP StorageWorks NAS 4000s or 9000s server. Overview information as well as procedural instructions are included in this guide.

NAS 4000s and 9000s Administration Guide
First Edition (November 2003)
Part Number: 352405-001

Contents

About This Guide

This administration guide provides information to help administrators:

- Plan the storage configuration
- Set up physical storage
- Manage users and groups
- Manage folders and shares
- Manage a UNIX® file system
- Manage a NetWare file system
- Remotely access the NAS server

“About this Guide” topics include:

Overview

This section covers the following topics:

Intended Audience

This book is intended for use by system administrators who are experienced with setting up and managing a network server.

Prerequisites

Before beginning, make sure you consider the items below.

- Knowledge of Microsoft® Windows® NT® or Windows Storage Server 2003 operating system
- Knowledge of HP hardware
- Location of all documentation shipped with your server

Conventions

Conventions consist of the following:

- Document Conventions
- Text Symbols
- Equipment Symbols

Document Conventions

The document conventions included in Table 1 apply in most cases.

Table 1: Document Conventions

Element	Convention
Cross-reference links	Figure 1
Key and field names, menu items, buttons, and dialog box titles	**Bold**
File names, application names, and text emphasis	Italics
User input, command and directory names, and system responses (output and messages)	`Monospace font` `COMMAND NAMES` are uppercase monospace font unless they are case sensitive
Variables	*`<monospace, italic font>`*
Website addresses	Sans serif font text: http://www.hp.com

Text Symbols

The following symbols may be found in the text of this guide. They have the following meanings:

WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or death.

Caution: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or data.

Note: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Equipment Symbols

The following equipment symbols may be found on hardware for which this guide pertains. They have the following meanings:

Any enclosed surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts.

WARNING: To reduce the risk of personal injury from electrical shock hazards, do not open this enclosure.

Any RJ-45 receptacle marked with these symbols indicates a network interface connection.

WARNING: To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.

Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. Contact with this surface could result in injury.

WARNING: To reduce the risk of personal injury from a hot component, allow the surface to cool before touching.

Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

WARNING: To reduce the risk of personal injury from electrical shock, remove all power cords to completely disconnect power from the power supplies and systems.

Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

WARNING: To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

Rack Stability

Rack stability protects personnel and equipment.

WARNING: To reduce the risk of personal injury or damage to the equipment, be sure that:

- The leveling jacks are extended to the floor.
- The full weight of the rack rests on the leveling jacks.
- In single rack installations, the stabilizing feet are attached to the rack.
- In multiple rack installations, the racks are coupled.
- Only one rack component is extended at any time. A rack may become unstable if more than one rack component is extended for any reason.

Getting Help

If you still have a question after reading this guide, contact an HP authorized service provider or access our website: http://www.hp.com.

HP Technical Support

Telephone numbers for worldwide technical support are listed on the following HP website: http://www.hp.com/support/. From this website, select the country of origin.

Note: For continuous quality improvement, calls may be recorded or monitored.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

HP Storage Website

The HP website has the latest information on this product, as well as the latest drivers. Access storage at: http://www.hp.com/country/us/eng/prodserv/storage.html. From this website, select the appropriate product or solution.

HP Authorized Reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-345-1518
- In Canada, call 1-800-263-5868
- Elsewhere, see the HP website for locations and telephone numbers: http://www.hp.com.

System Overview

1

The HP StorageWorks NAS server can be used in many types of computing environments, from basic Microsoft Windows workgroups to complicated multiprotocol domains using DFS, NFS, FTP, HTTP, and Microsoft SMB. The corresponding varieties of clients that can be serviced include any Windows, UNIX, Linux, Novell, or Macintosh variant.

This chapter provides an overview of these environments and deployments and includes brief descriptions of system user interfaces, applications, and options.

Product Definition and Information

The business class NAS 4000s and the enterprise class NAS 9000s are solutions that provide reliable performance, manageability, and continuous data availability through the fusion of Network Attached Storage (NAS) and Storage Area Network (SAN) technologies.

Server Hardware and Software Features

Refer to the *HP StorageWorks NAS 4000s/9000s Installation Guide* for a listing of server hardware and software features.

For specific software product recommendations, go to the HP website:

http://www.hp.com/go/nas

Product Manageability

The NAS server ships with the following utilities and features that ease the administration tasks associated with managing the system:

- The Rapid Startup Wizard is a user friendly configuration utility that ensures easy configuration.
- The WebUI is a simple, graphical user interface (GUI) that helps with administration tasks.

- Insight Manager is a comprehensive tool designed to be a key component in the systems management environment. It monitors the operations of HP servers, workstations, and clients. Insight Manager provides system administrators more control through comprehensive fault and configuration management, and industry leading remote management.
- The Integrated Lights-Out feature provides remote access, sends alerts, and performs other management functions, even if the operating system of the host server is not responding.

Product Redundancy

The NAS server is specifically designed to perform file serving tasks for networks. Using industry standard components, redundancy of power supplies, NICs, and fans ensures reliability.

The clustering ability of the NAS device further ensures continuous data availability, because data being processed by one server transitions over to the other server in a failover situation.

Other industry standard features, such as redundant array of independent drives (RAID) and remote manageability, further enhance the overall dependability of the NAS server.

The server contains dual 36.4 GB hard drives preconfigured with the NAS operating system so that the active system volume is mirrored (RAID 1+0) to the second drive. If one of the internal drives fails, the integrity of the system is preserved, because the system will use the copy of the operating system on the remaining healthy drive. The drives in the server are hot-pluggable, so the failed drive can be replaced while the system is running. When the failed drive is replaced, the system automatically uses the version of the operating system on the healthy drive to rebuild the replacement.

The NAS server includes dual power supplies. A power supply can be replaced while the server is running. To ensure redundancy, it is important to connect each power supply to a separate power source. If one power source fails, the server remains operational through the second power source.

Through a seamless, hardware-based, graphical remote console, the Integrated Lights-Out port provides the administrator with full control of the server from a remote location. Using a client browser, the administrator can remotely power up, power down, and operate the console. A built in processor makes the port independent of the server and the operating system.

Deployment Scenarios

Various deployment scenarios are possible. Typical application of NAS devices include:

- **File server consolidation**

  As businesses continue to expand their information technology (IT) infrastructures, they must find ways to manage larger environments without a corresponding increase in IT staff. Consolidating many servers into a single NAS device decreases the number of points of administration and increases the availability and flexibility of storage space.

- **Multiprotocol environments**

  Some businesses require several types of computing systems to accomplish various tasks. The multiprotocol support of the NAS server allows it to support many types of client computers concurrently.

- **Protocol and platform transitions**

  When a transition between platforms is being planned, the ability of the NAS server to support most file sharing protocols allows companies to continue to invest in file storage space without concerns about obsolescence. For example, an administrator planning a future transition from Windows to Linux can deploy the NAS server with confidence that it can support both CIFS and NFS simultaneously, assuring not only a smooth transition, but also a firm protection of their investment.

- **Remote office deployment**

  Frequently, branch offices and other remote locations lack dedicated IT staff members. An administrator located in a central location can use the WebUI of the NAS server, Microsoft Terminal Services, and other remote administration methods to configure and administer all aspects of the NAS server.

Configuration Options for the NAS Server

The NAS server can be deployed in two different configurations attached to a storage area network (SAN): as a single NAS device or as a clustered pair. By default, the NAS server does not ship with a standard HBA to allow for maximum flexibility in the SAN environment. For a list of supported HBAs for the NAS servers, see the quick specs for the product.

NAS Server as a Single Device

In the single NAS device configuration, a NAS server is attached to a SAN via a single or pair of fiber channel host bus adapters (HBAs) and one or more SAN switches. In Figure 1, the dual HBA approach is illustrated. Each HBA should be connected to a separate switch that has access to the same controller pairs. This connection method allows redundant paths to the same storage subsystem. Dual HBAs per NAS device is recommended but not required for stand alone deployments, since dual HBAs allow for path failure while still providing access to the data.

SAN storage is not managed by the NAS server and requires coordination between the NAS administrator and the SAN administrator. LUN requests need to be made and assigned to the NAS server using selective storage presentation from the SAN controller pairs. Naming of the LUNs and the connections to the NAS server is important for tracking and identification purposes. At least one LUN is required by the NAS server for the device to be functional.

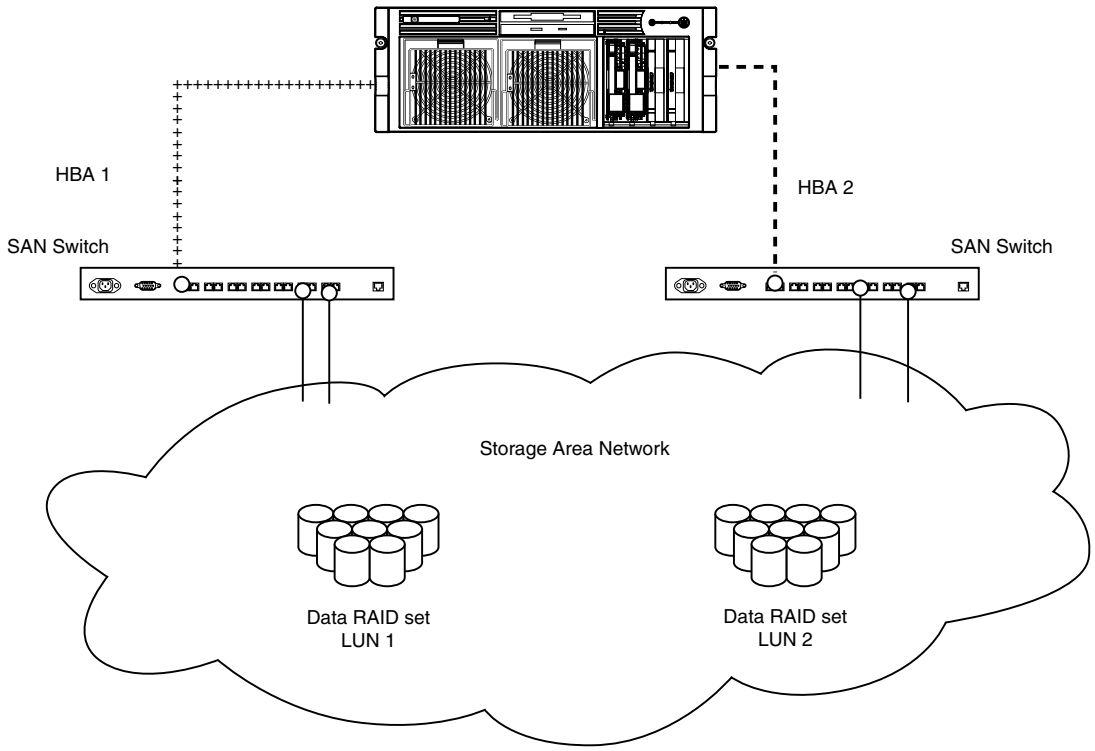

Figure 1: NAS server as a single device

NAS Server as a Clustered Pair

In the clustered configuration, two NAS devices are linked via a private network and have access to shared storage as illustrated in Figure 2. In clustered deployments, it is recommended that each NAS device be attached to the SAN via a pair of fiber channel HBAs. Dual HBAs per NAS device is recommended but not required for cluster deployments, since dual HBAs allow for path failure without causing node failover in the cluster.

For the NAS server, cluster setup requires at least three LUNs. One LUN is required for the quorum disk and two LUNs are required for data disks. Individual data disks specific to each node are required for the proper setup of file sharing where both nodes participate in file-sharing activities. Clustered NAS systems have the ability to provide redundant active/active access to file shares on disk. However, as with all Microsoft-based clusters, the unit of ownership among nodes is at the disk level. Therefore, individual fileshares can be accessed by only one node at a time based on which node owns the disk.

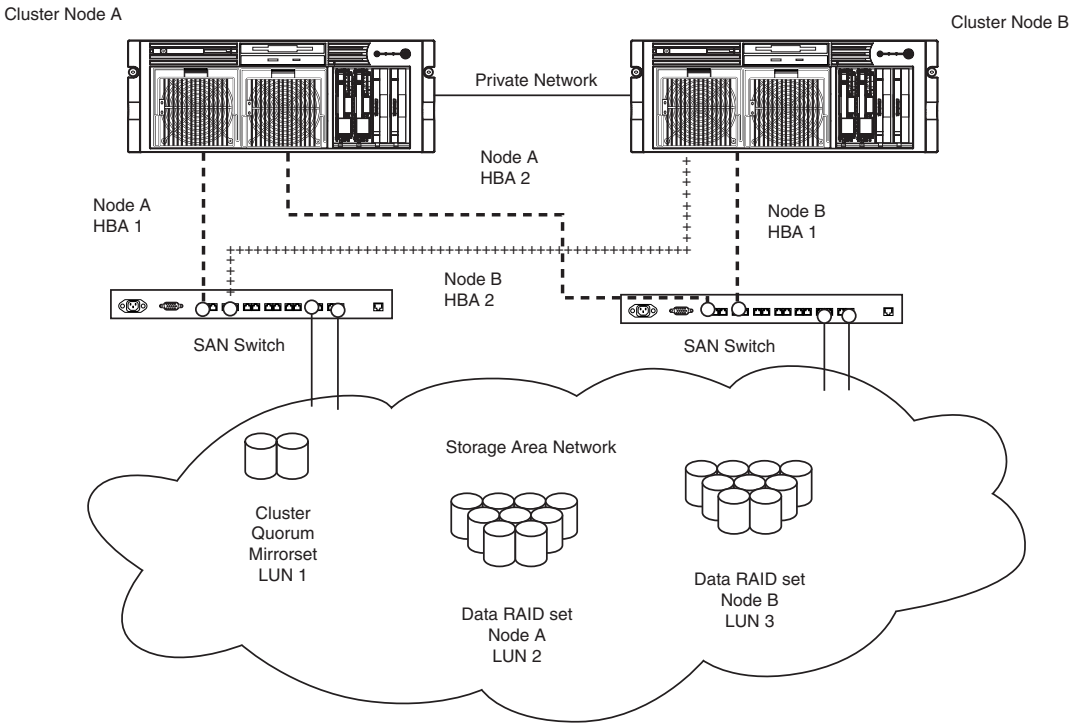

Figure 2: NAS server as a clustered pair of devices

Multi Node Support Beyond Two Nodes

The NAS 4000s and 9000s devices may be deployed in multi node clustering beyond two nodes. Refer to the associated Storage Array documentation to determine the number of nodes supported by the array under Windows Storage Server 2003. While the discussion presented in this guide addresses only two nodes, additional nodes may be added into the cluster. Considerations for additional fiber path connections and the private network should be made. In the case of the private network, a hub or switch is required since the cross over cable is no longer applicable.

Connecting NAS Servers to the Network

The NAS server can be connected to the network as a single device or as a clustered pair.

NAS Server Single Device Deployment

In a single device deployment, network connections are divided into two categories: client data segments and management segments. The default shipping configuration contains a two-port network interface controller (NIC) 10/100/100 that provides one port for management and one port for client data. The management port is intended for use with the 3202 port of the device to enable use of the WebUI that accompanies the product. It is from this WebUI that most management and administrative procedures can be accomplished. An additional management port for remote console and diagnostics is provided off the Integrated Lights-Out (iLO) port. HP recommends that this connection be placed on a management LAN separate from the corporate infrastructure.

The NAS server supports the use of NIC teaming. NIC teaming provides failover and load balancing of network ports of the NAS server. NIC teaming requires the network cables to be installed on the same subnet to enable it to work. For example in Figure 3, Eth1 and Eth2 would need to be on the same subnet to enable this functionality on those ports. However, it is not recommended to assign IP addresses to the ports that will be teamed or load balanced prior to the installation and setup of NIC teaming. For this reason, HP recommends that you set all network ports to DHCP. For information concerning the configuration of NIC teaming after setup is complete, see Chapter 2.

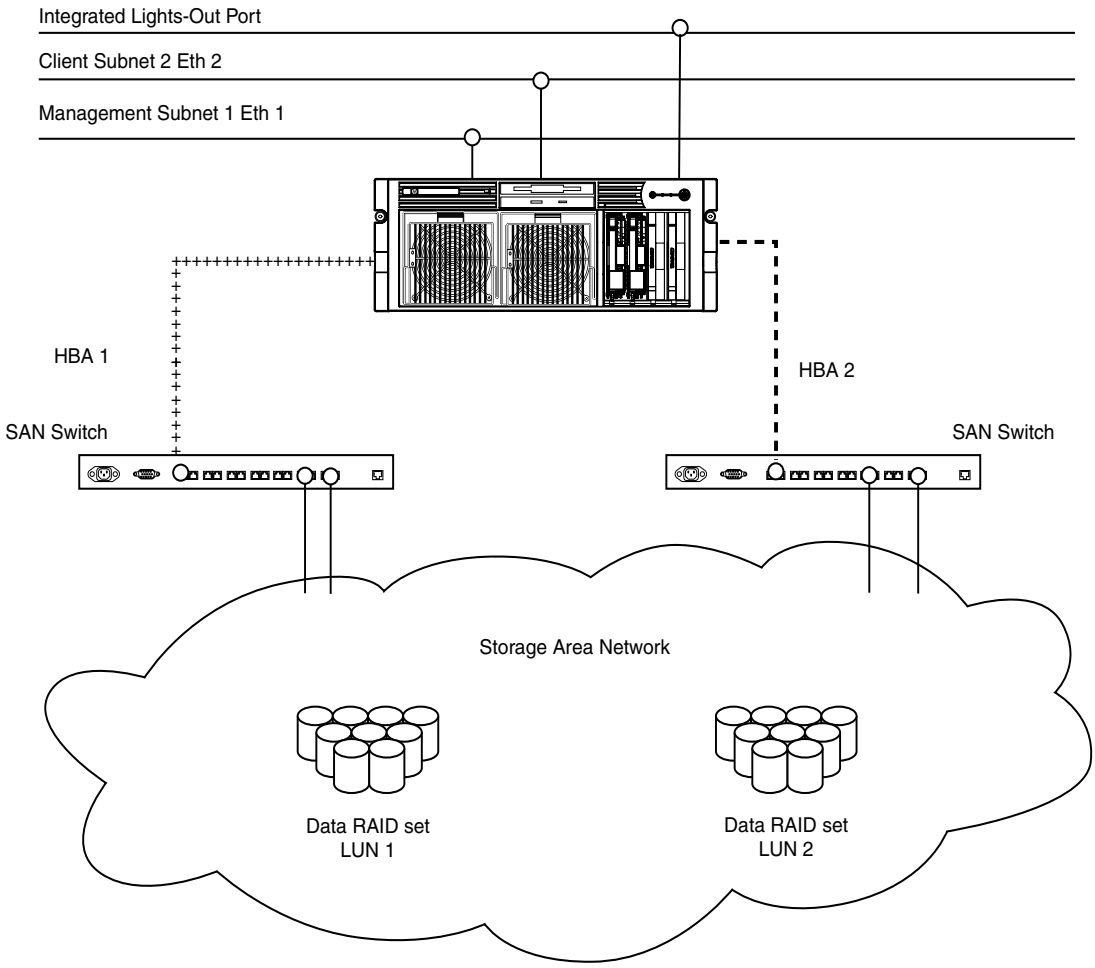

Figure 3: NAS server single device deployment

NAS Server Cluster Deployment

In a clustered deployment, network connections are divided into three categories: cluster maintenance segments, client data segments, and management segments. The default shipping configuration contains a two-port 10/100/1000 NIC that together provide one port for the cluster maintenance and one port for management and client data. An additional management port for remote console and diagnostics is provided off of the iLO port. HP recommends that this connection be placed on a management LAN separate from the corporate infrastructure.

The cluster maintenance segment is also known as the heartbeat connection between the nodes. In standard configurations, this connectivity is obtained by the use of a crossover network cable. The same effect can be achieved by placing the heartbeat ports on its own switch or VLAN, as illustrated in Figure 4 as Eth1. The purpose is to isolate and guarantee the connectivity between the nodes without interruption. If interruption occurs, the remaining cluster node will assume the other node has gone down and initiate a failover. A second cluster heartbeat path is often recommended as a redundant path. The redundant path is often done over one of the remaining network segments and is not dedicated.

The client data segments of a cluster must reside on identical network segments for each node. As illustrated in Figure 4, Eth2 from both nodes is shown on the same segment. This co-location on the same segment allows cluster resources such as file shares to failover to the second node and continues to provide availability to the resources. The failed resource will have the same IP address after the failover and therefore must have a port corresponding to the identified subnet of the IP address present on both nodes.

The NAS server supports the use of NIC teaming in clustered configurations. All previous caveats from the “NAS server single node deployment” on networks section still apply.

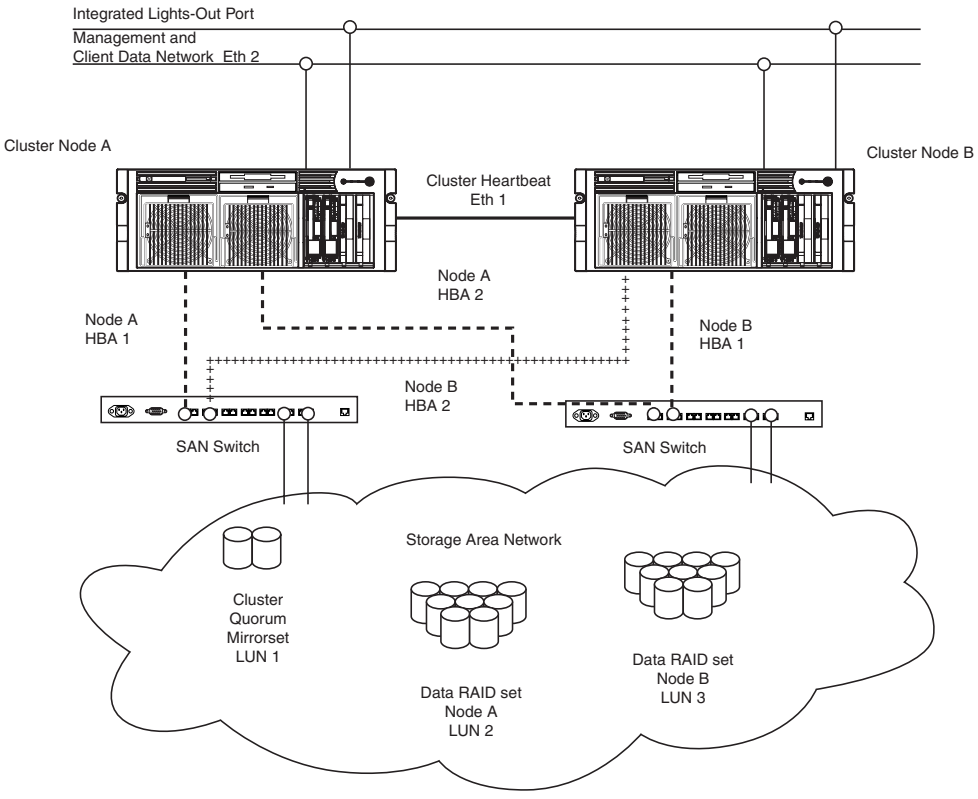

Figure 4: NAS server cluster deployment

Environment Scenarios

The NAS server is deployed in one of two security modes:

- Workgroup
- Domain (Windows NT Domain or Active Directory Domain)

The NAS server uses standard Windows user and group administration methods in each of these environments. For procedural instructions on managing users and groups, see Chapter 6 of this guide.

Regardless of the deployment, the NAS server integrates easily into multiprotocol environments, supporting a wide variety of clients. The following protocols are supported:

- Distributed File System (DFS)
- Network File System (NFS)
- Hypertext Transfer Protocol (HTTP)
- File Transfer Protocol (FTP)
- Microsoft Server Message Block (SMB)

Workgroup

In a workgroup environment, users and groups are stored and managed separately, on each member server of the workgroup. Workgroups are typical for very small deployments where little or no computing environment planning is required.

Note: In a clustered deployment, the clusters must be members of a domain. Therefore, workgroup environments are supported only in non-clustered deployments.

Domain

When operating in a Windows NT or Active Directory domain environment, the NAS server is a member of the domain and the domain controller is the repository of all account information. Client machines are also members of the domain and users log on to the domain through their Windows based client machines. The domain controller also administers user accounts and appropriate access levels to resources that are a part of the domain. Additional information about planning for domain environments can be found at:

http://www.microsoft.com/windowsserver2003/technologies/directory/activedirectory/default.mspx

The NAS server obtains user account information from the domain controller when deployed in a domain environment. The NAS server itself cannot act as a domain controller, backup domain controller, or the root of an Active Directory tree as these functions are disabled in the operating system.

User Interfaces

There are several user interfaces that administrators can use to access and manage the NAS server. Two of these interfaces are:

- NAS server WebUI
- NAS server Desktop

Each interface contains the same or similar capabilities, but presents them in a different manner. Each of these interfaces are illustrated in the following sections.

NAS Server Web-Based User Interface

The WebUI provides for system administration, including user and group management, share management, and local storage management.

Refer to the *HP StorageWorks NAS 4000s/9000s Installation Guide* for detailed information on using the Rapid Startup Wizard for initial setup.

To access the WebUI, launch a Web browser and enter the following in the address field:

```
https://<your NAS machine name or IP Address>:3202/
```

The default user name is `Administrator`. The default password is `hpinvent`. Online help for the WebUI is available by clicking the **Help** tab on the primary WebUI screen.

The primary screen of the WebUI is shown in Figure 5.

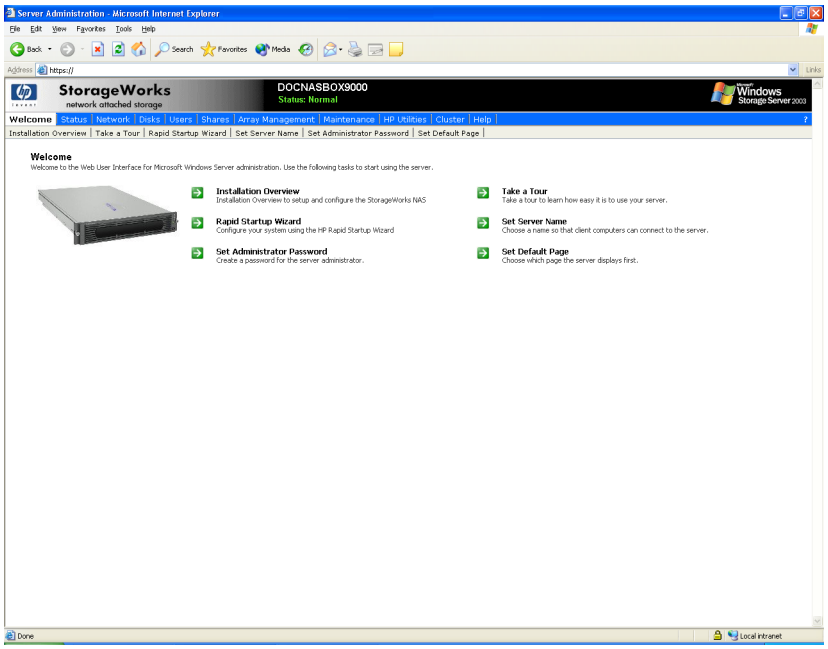

Figure 5: Primary WebUI screen

As shown in Figure 5, the following areas are administered through this interface:

Menu Tabs

Status

The Status option displays alerts generated by the WebUI.

Network

The Network option contains system settings, including system identification, global settings, interfaces settings, administration settings, Telnet settings, and SNMP settings.

Disks

Use this option to manage disks, volumes, disk quotas, and shadow copies.

Users

Use this option to manage local users and groups.

Shares

The administrator creates folders and shares to control access to files. When a share is created, the administrator indicates the protocols that can be supported by that share as well as the users and groups of users that have access. Protocol parameters are entered in this Shares option. See Chapter 6 for additional information.

Array Management

Manage arrays and pathing software from this tab.

Maintenance

Maintenance tasks include setting date and time, performing system restarts and shutdowns, viewing audit logs, setting up Email alerts, linking to remote management, and selecting and configuring your UPS.

HP Utilities

Access HP system management utilities such as remote management, enable floppy boot, File and Print Services for NetWare, and the HP System Management WebUI.

Cluster

Use this option to configure and manage the cluster.

Help

This option contains help information for the WebUI.

Welcome Screen Contents

Installation Overview

Use to set up and configure the NAS server. This is an online, supplemental guide. A more comprehensive paper document is provided in the country kit that shipped with the server.

Rapid Startup Wizard

Use this utility to enter system setup and configuration information.

Set Administrator Password

Create a password for the server appliance administrator.

Take a Tour

Learn how to use the NAS server.

Set Server Name

Choose a name so that client computers can connect to the server.

Set Default Page

Choose which page the server appliance displays first.

NAS Server Desktop

The NAS server desktop can be accessed by:

- Directly connecting a keyboard, mouse, and monitor
- Using the WebUI Maintenance tab and selecting **Remote Desktop**
- Using the Integrated Lights-Out port

Note: When using Remote Desktop to connect to the NAS desktop do not use the window close feature ( ). Click on **Start/Log Off Administrator** to exit Remote Desktop. See "Improper Closure of Remote Desktop" in Chapter 2.

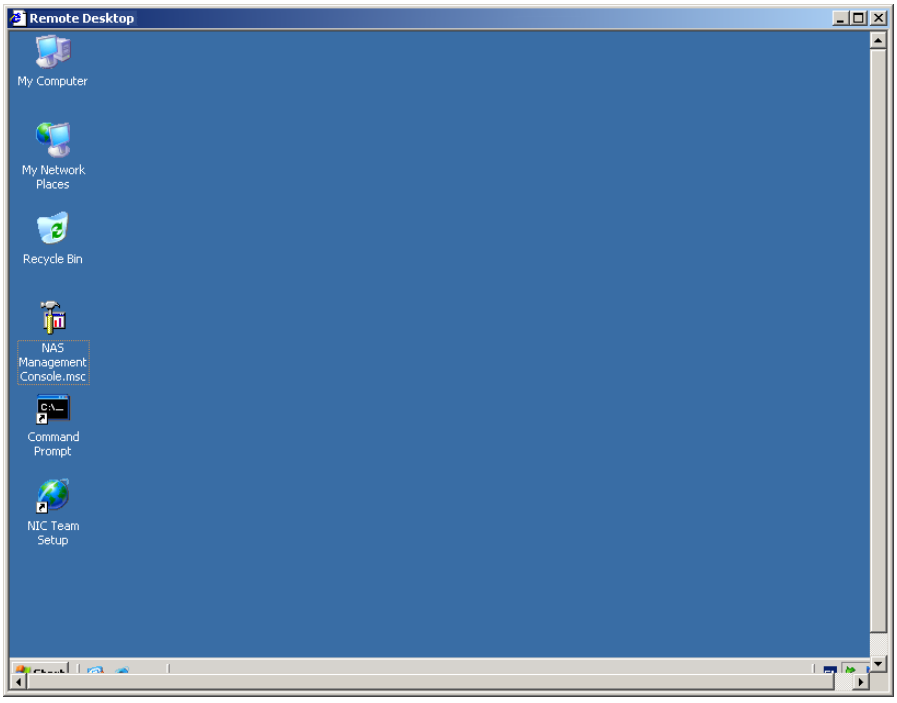

Figure 6: NAS server desktop

The following icons are available from the Desktop:

- NAS Management Console
- NIC Team Setup

NAS Management Console

Click this icon to access the following folders:

- **Core Operating System** is used to manage local users and groups, access performance logs and alerts, and manage the event viewer.
- **Disk System** contains access to the HP Array Configuration Utility and local disk management, including a volume list and a graphical view of the disks.
- **File Sharing** contains modules for the configuration of file sharing exports. CIFS/SMB (Windows) and NFS (UNIX) file shares are managed through this folder.
- **System** contains system summary information.

NIC Team Setup

Click this icon to install the HP Network Teaming and Configuration utility. See Chapter 2 for additional information on this feature.

Basic Administrative Procedures and Setup Completion

2

Basic system administration functions are discussed in this chapter.

This chapter also continues the process of setting up the system that was started using the *HP StorageWorks NAS 4000s/9000s Installation Guide* by discussing additional setup procedures and options. Further steps can also be viewed online by clicking the **Installation Overview** tab from the primary WebUI screen.

Unless otherwise instructed, all procedures are performed using the NAS Web Based User Interface (WebUI).

Note: The NAS Desktop can be accessed via a directly connected keyboard, mouse, and monitor through Remote Desktop, or by using an Integrated Lights-Out port.

Basic Administrative Procedures

Basic administrative procedures include:

- Setting the system date and time
- Shutting down or restarting the server
- Viewing and maintaining audit logs
- Using Remote Desktop
- Setting up e-mail alerts
- Changing system network settings

These functions are performed in the **Maintenance** menu of the WebUI.

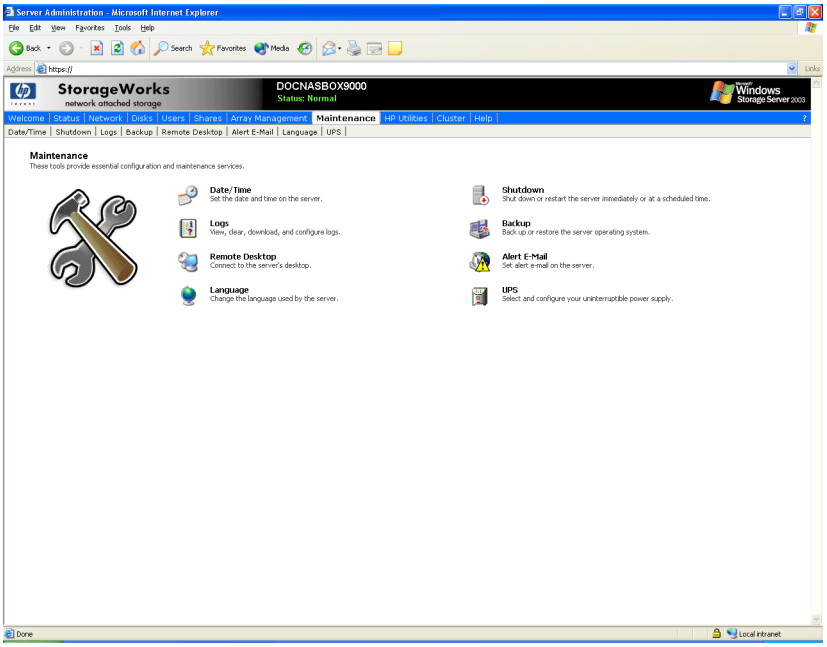

Figure 7: Maintenance menu

Setting the System Date and Time

To change the system date or time:

1. From the WebUI, select **Maintenance** and **Date/Time**. The **Date and Time Settings** dialog box is displayed.
2. Enter the new values and then click **OK**. The **Maintenance** menu is displayed.

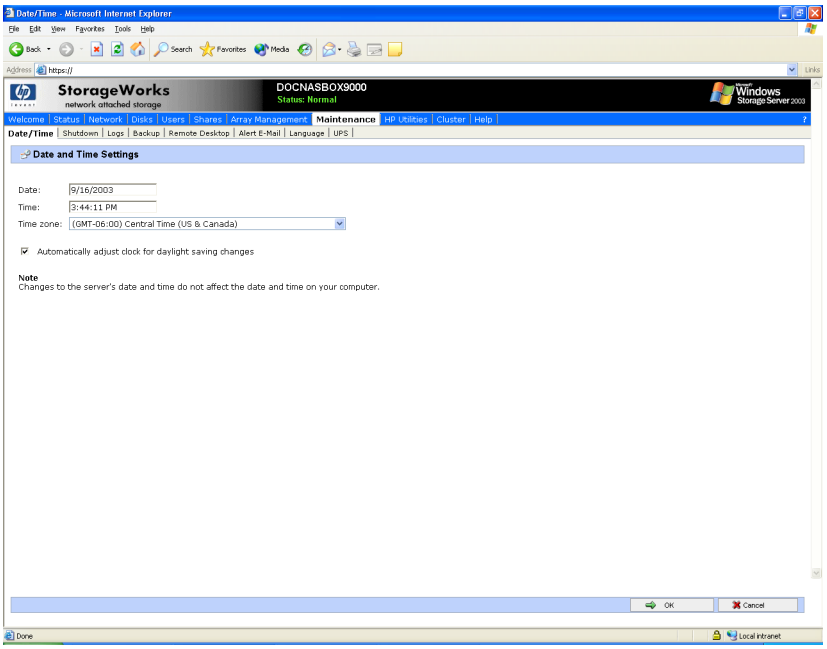

Figure 8: Date and Time dialog box

Shutting Down or Restarting the Server

Caution: Notify users before powering down the system. Both UNIX and Windows NT users can be drastically affected if they are not prepared for a system power-down.

1. From the NAS server WebUI, select **Maintenance**, **Shutdown**. Several options are displayed: **Restart**, **Shut Down**, and **Scheduled Shutdown**.

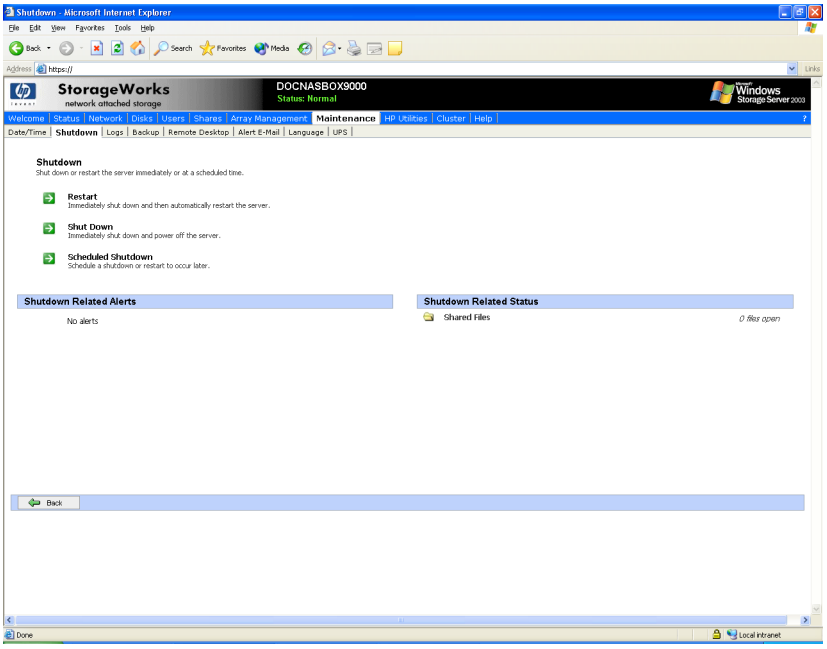

Figure 9: Shutdown menu

   a. To shut down and automatically restart the server, click **Restart**.
   b. To shut down and power off the server, click **Shut Down**.
   c. To schedule a shutdown, click **Scheduled Shutdown**.
2. Regardless of the choice, a confirmation prompt is displayed. After verifying that this is the desired action, click **OK**.

Note: Client computers connected to the server will receive a warning message prior to shutdown. Clients connected via Remote Desktop do not receive any warning.

Viewing and Maintaining Audit Logs

A variety of audit logs are provided on the NAS server. System events are grouped into similar categories, representing the seven different logs.

To access the logs from the WebUI, select **Maintenance**, **Logs**. The **Logs** menu is displayed.

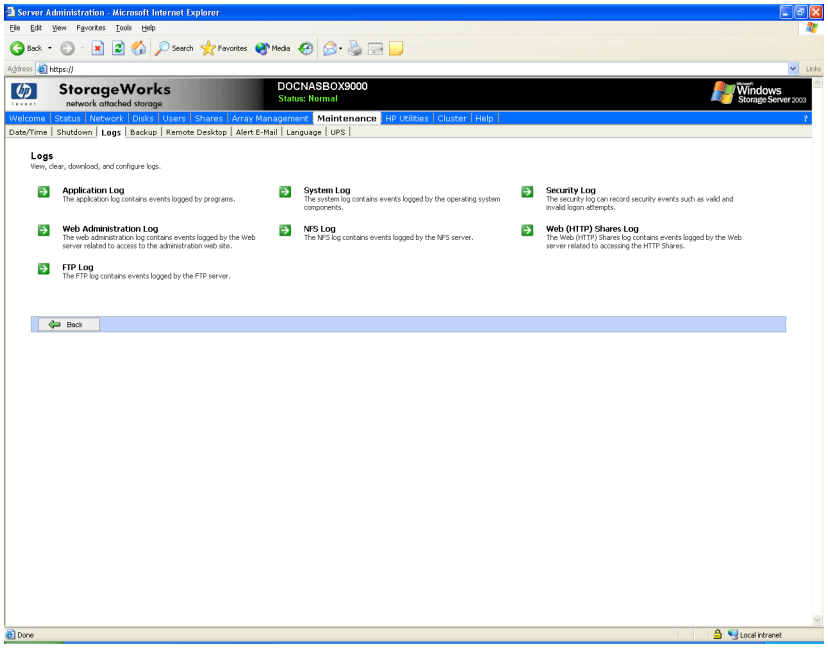

Figure 10: Logs menu

A variety of logs are available and are listed in Figure 10.

Each log has viewing, clearing, printing, and saving options.

Note: You should not use the WebUI to view log files greater than 2 MB. Select Log properties to adjust the maximum file size, or download the file to view.

Note: NFS logging is disabled by default. Enable NFS logging using the NAS Management Console. NFS stops logging when the log file is full.

Using Remote Desktop

Remote Desktop is provided in the WebUI to allow for additional remote system administration and the use of approved third-party applications. Backup software and antivirus programs are examples of approved applications.

In addition, Remote Desktop is used to access the NAS Management Console of the NAS device.

To open a Remote Desktop session from the WebUI, select **Maintenance**, **Remote Desktop**. A Remote Desktop session is opened. Enter the appropriate password to log on to the server.

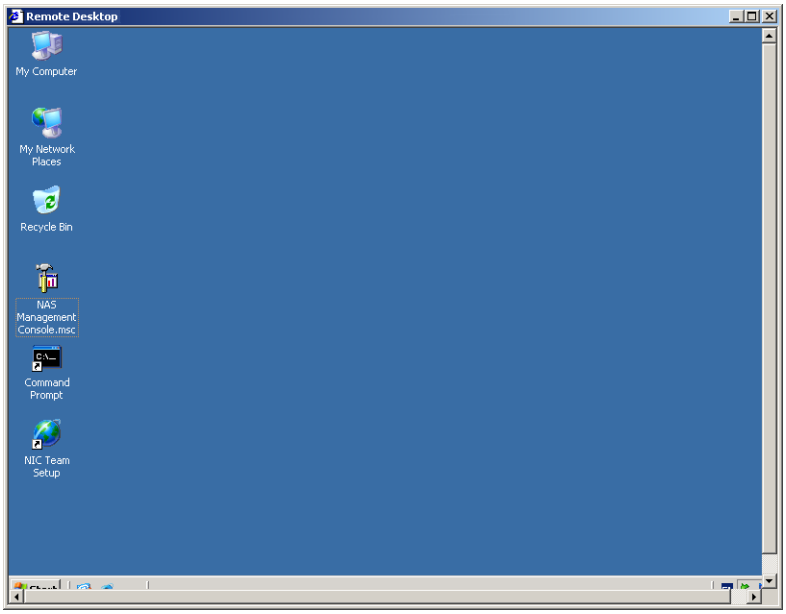

Figure 11: Remote Desktop session

Caution: Two open sessions of Remote Desktop are allowed to operate at the same time. After completing an application do not use the window close feature ( ) to close that session of Remote Desktop. Click **Start/Log Off Administrator** to exit Remote Desktop.

Improper Closure of Remote Desktop

Certain operations can leave the utilities running if the browser is closed versus exiting from the program via the application menu or logging off the Remote Desktop session. The default timeout has been set to 15 minutes but it may require up to 30 minutes for the application to exit. This value may be adjusted in the **Terminal Services Configuration** option under Administrator Tools, accessed either through the desktop or via a Remote Desktop session.

A maximum of two Remote Desktop sessions may be used at any given time. Improper exit from a session can result in the sessions becoming consumed. Sessions and processes can be terminated using the **Terminal Services Manager** via **Start >Programs >Administrator Tools**.

Note: The Terminal Services Manager must be accessed via the iLO port or direct attached console.

Setting up E-mail Alerts

E-mail messages are limited to the alerts generated from the WebUI status bar or the WebUI status page, as well as some event log messages. Some alerts, such as the restart of the server, only occur if the WebUI was utilized to initiate the request. For example, a restart initiated using the WebUI will general an e-mail message indicating a restart has occurred. Initiating a restart using the Windows Storage Server 2003 schedule or Desktop will not. Messaging in the status bar and page is limited to the following areas:

- WebUI Alerts
  - NTBackup backup started
  - NTBackup restore started
  - Defrag started
  - UPS power failure
  - Restart pending
  - Shutdown pending
  - DFS not configured
  - Date and time not configured
  - No certificate
  - Quota management alerts
- Event Log Messages
  - NTBackup Information
  - UPS power failed
  - UPS power restored
  - UPS invalid config
  - UPS system shutdown
  - Quota management alerts

To activate this option:

1. From the WebUI, select **Maintenance**, **Alert E-mail**. The **Set Alert E-Mail** dialog box is displayed.
2. Select **Enable Alert E-mail**.

3. Indicate the types of messages to be sent.
   - Critical alerts
   - Warning alerts
   - Informational alerts
4. Enter the desired e-mail address in the appropriate boxes.
5. After all settings have been entered, click **OK**.

Changing System Network Settings

Network properties are entered and managed from the **Network** menu. Most of these settings are entered as part of the Rapid Startup process. Settings made from this menu include adding the NAS server to a domain.

Online help is available for these settings. Figure 12 is an illustration of the Network settings menu.

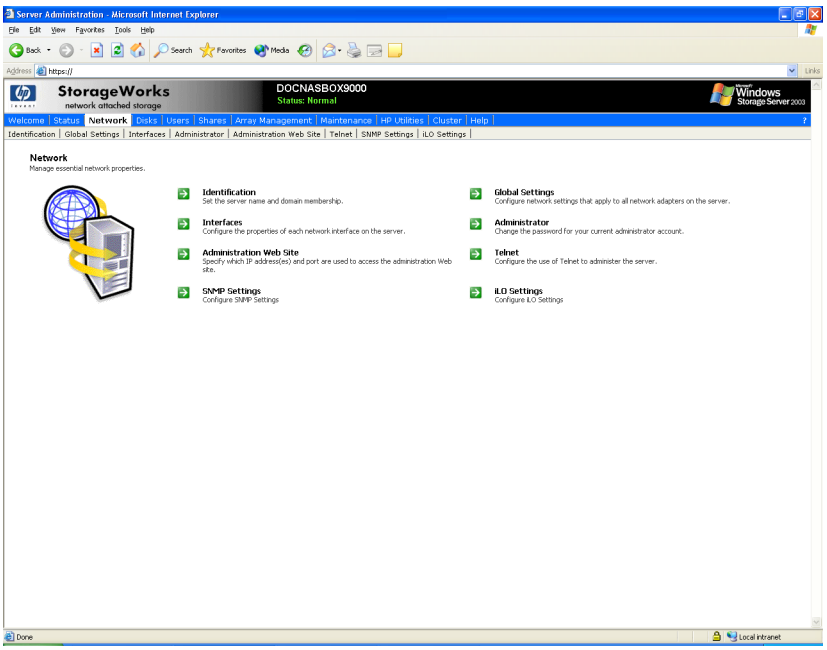

Figure 12: Network menu

Setup Completion

After the NAS device is physically set up and the basic configuration is established, additional setup steps must be completed. Depending on the deployment scenario of the NAS device, these steps may vary.

Additional setup steps may include:

- Activating the iLO port using the license key
- Setting up Ethernet NIC teams (optional)
- Managing system storage
- Creating and managing users and groups
- Creating and managing file shares

Each of these setup steps is discussed in the following sections.

Activating the iLO Port Using the License Key

The Remote Desktop feature of the iLO port requires a license key. The key is included with the product inside the Country Kit. Refer to the iLO Advanced License Pack for activation instructions.

To access the iLO port, click on **HP Utilities**, then click **Remote Management**.

Setting up Ethernet NIC Teams (Optional)

The NAS server is equipped with the HP Network Teaming and Configuration utility. The utility allows administrators to configure and monitor Ethernet network interface controllers (NIC) teams in a Windows-based operating system. These teams provide options for increasing fault tolerance and throughput.

Fault tolerance provides automatic redundancy. If the primary NIC fails, the secondary NIC takes over. Load Balancing provides the ability to balance transmissions across NICs.

Note: The NAS server ships with the NIC teaming utility available, however it must be installed and configured.

Note: Installing NIC teaming requires a restart of the server.

Procedures include:

- Installing the HP Network Teaming utility
- Opening the HP Network Teaming utility
- Adding and configuring NICs in a team
- Configuring the NIC team properties
- Checking the status of the team
- NIC teaming troubleshooting

Installing the HP Network Teaming Utility

Before using the HP Network Teaming utility, it must be installed.

Note: Installing and configuring NIC teaming should always be performed via iLO port or the console using a direct attached keyboard, monitor, and mouse since IP connections could be reset during the configuration process. Do not use Remote Desktop.

To install the HP Network Teaming utility:

1. In the URL field of the Web browser, enter the IP address of the Integrated Lights-Out port.

Note: The iLO port requires a license key. The key is included with the product inside the Country Kit. Refer to the iLO Advanced License Pack for activation instructions.

Note: The iLO port can also be accessed from the HP Utilities tab of the WebUI by clicking the remote management link.

2. At the Integrated Lights-Out Account Login screen, supply the username and password for the iLO and click **Login**.
3. Click the Remote Console tab. The Remote Console Information screen is displayed.
4. Click on the Remote Console choice in the menu on the left side of the screen.
5. Press the **Ctrl-Alt-Del** button to login to the console.
6. Supply an administrator username and password. The NAS server desktop is displayed.
7. Double-click the **NIC Team Setup** icon on the desktop.
8. When the following message box is displayed, click **Install**.

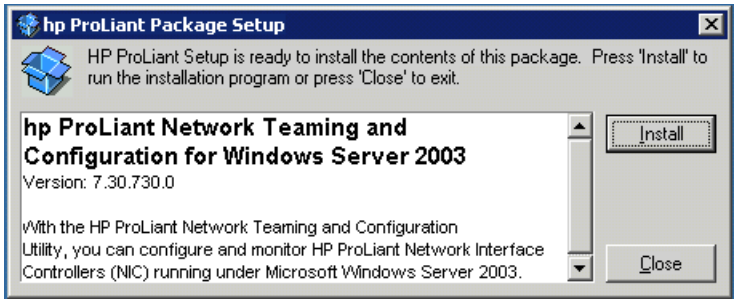

Figure 13: Installing Network Teaming

9. When the installation process is complete, the following screen is displayed. Click **Close**.

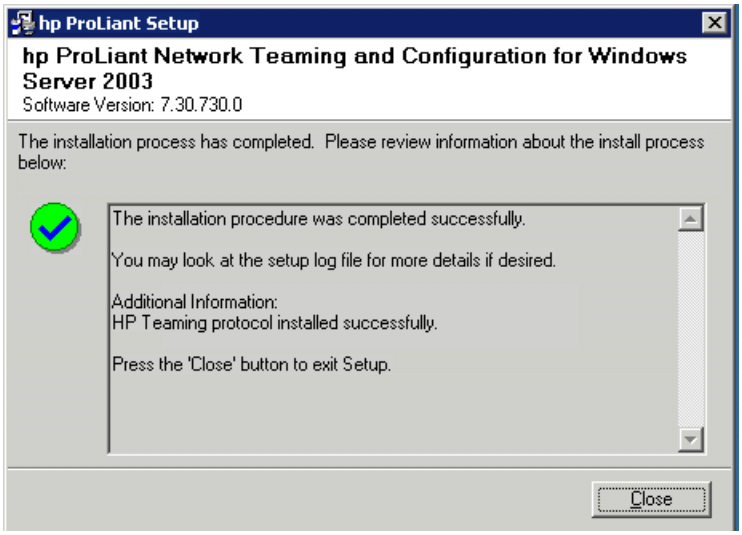

Figure 14: Network Teaming installation complete

10. Restart the system.

Caution: To ensure proper functioning of the software, the server must be restarted at this time.

Opening the HP Network Teaming Utility

The HP Network Teaming utility is now accessible from the Windows toolbar at the bottom of the NAS server desktop. To open the utility, click the **HP Network Teaming utility** icon.

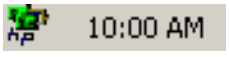

Figure 15: HP Network Teaming utility icon

Adding and Configuring NICs in a Team

Before a NIC is teamed, verify the following:

- The NICs must be on the same network.
- The NICs must be DHCP enabled and the DNS server address must be left blank.

Note: The teaming utility becomes unstable if static IP addresses, subnets, and DNS addresses are set before teaming.

- Duplex and speed settings must be set to use the default values.

To team the NICs:

1. Open the HP Network Teaming utility. The **Network Teaming and Configuration Properties** dialog box is displayed. The type of NIC and the slot and port used is shown.

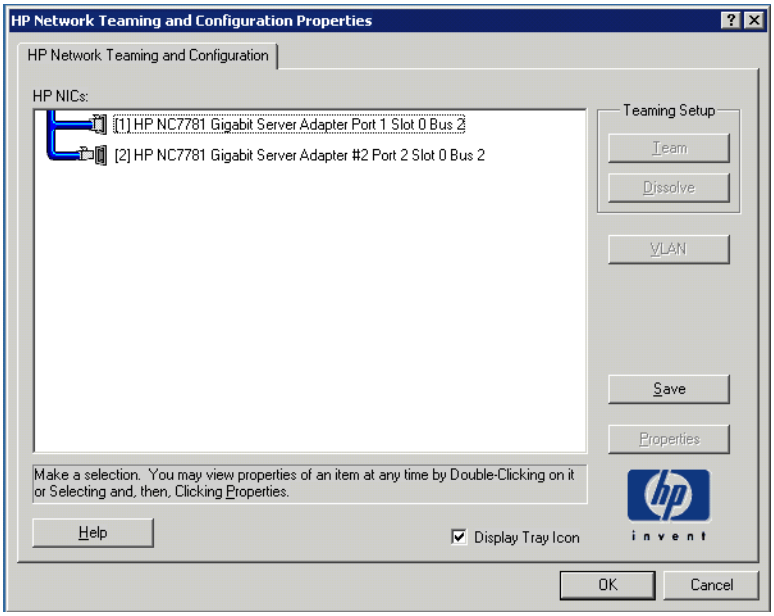

Figure 16: HP Network Teaming Properties dialog box

2. Highlight the NICs to team.
3. Click the **Team** button. The **Teaming Controls** tab of the Properties dialog box is displayed.

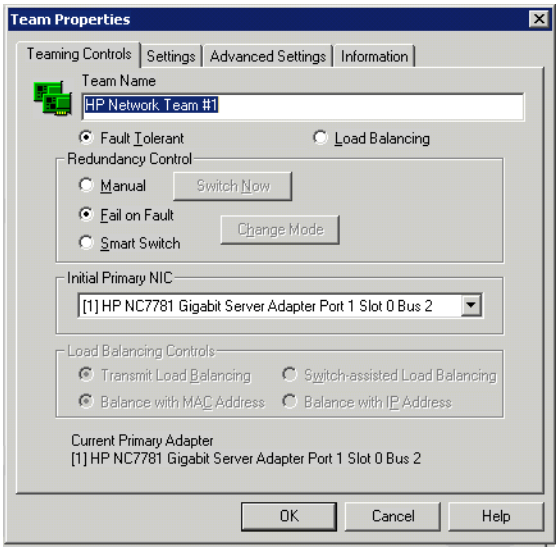

Figure 17: NIC Properties, Teaming Controls tab, Fault Tolerant option

4. Configure the team by choosing either **Fault Tolerant** or **Load Balancing**.

   The fault tolerance and load balancing options are discussed in the following sections.

Fault Tolerance

The Fault Tolerance teaming option provides three redundancy control options:

- **Manual**—This setting allows change from a Primary NIC to a Secondary NIC only when **Switch Now** is clicked.

Note: The **Switch Now** option is disabled until **Manual** is selected and then **OK** is clicked.

- **Fail on Fault**—This setting automatically switches from a primary NIC to a secondary NIC when the primary NIC fails.
- **Smart Switch**—This setting lets a member of a team be selected as the preferred Primary Smart Switch NIC. As long as this NIC is operational, it is always the active NIC. If the NIC fails and it is eventually restored or replaced, it automatically resumes its status as the active NIC.

Note: **Smart Switch** is the recommended choice for fault tolerance.

Detailed information about configuring teams for fault tolerance can be found in the HP Network Teaming utility help.

Load Balancing

The **Load Balancing** teaming option provides four load balancing control options:

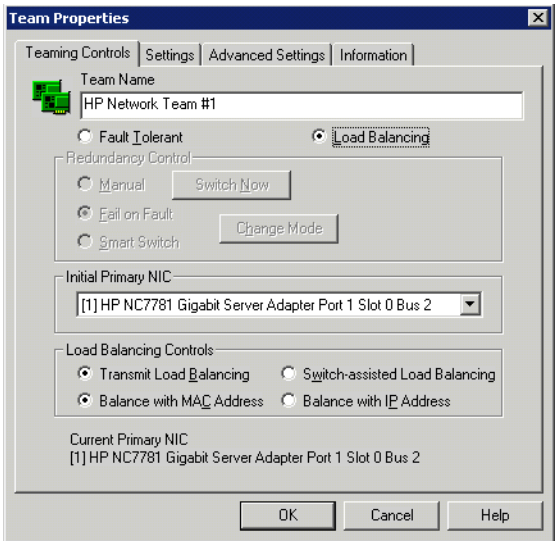

Figure 18: NIC Properties, Teaming Controls tab, Load Balancing option

Detailed information about these four load balancing teaming options can be found in the HP Network Teaming Help.

- **Transmit Load Balancing**—All transmit IP frames are load balanced among all team members based on a Load Balancing algorithm in the teaming device driver. The Current Primary adapter transmits all other frames, and receives all frames for the team. If a failover event occurs, one of the non-Primary adapters assumes the role of Current Primary adapter, and transmit IP packets are load balanced among all remaining team members. If a failure occurs in any of the non-Primary adapters, the packets are load balanced among all remaining team members.
- **Switch-assisted Load Balancing**—All transmit packets are load balanced among all team members based on a Load Balancing algorithm in the teaming device driver. The receive packets are load balanced among all team members by the switch. If a failure of any team member occurs, the packets are load balanced among the remaining adapters. There is no primary adapter in a Switch-assisted Load Balancing team.
- **Balance with MAC Address**—This feature allows load balancing of IP packets among the teamed NICs using the last four bits of the MAC Address. (See following Note.)
- **Balance with IP Address**—This feature allows load balancing of IP packets among the teamed NICs using the last four bits of the IP Address. (See following Note.)

Note: The teaming utility can load balance IP packets among the teamed NICs installed in a server. The primary NIC in the team receives all incoming packets. The choice is available to load balance with the source MAC address (the address transmitted from the workstation) or the source IP address.

Using the last four bits of either source address, the teaming driver algorithm assigns this source address to the port of one of the NICs in the team. This port is then used to transmit all packets destined for that source address. If there are four NICs in the team, the packets are received by the primary NIC on the team. The packets are retransmitted through one of the four ports.

---

5. Click **OK** to accept the team properties.

6. Click **OK** in the HP Network Teaming and Configuration Properties Screen to apply the changes.
7. Click **Yes** when prompted to apply all configuration changes. Wait while the adapters are configured. This process could take several seconds.
8. The following screen is displayed, indicating that there are additional procedures to perform in the NIC teaming process. Click **Yes** to reboot now.

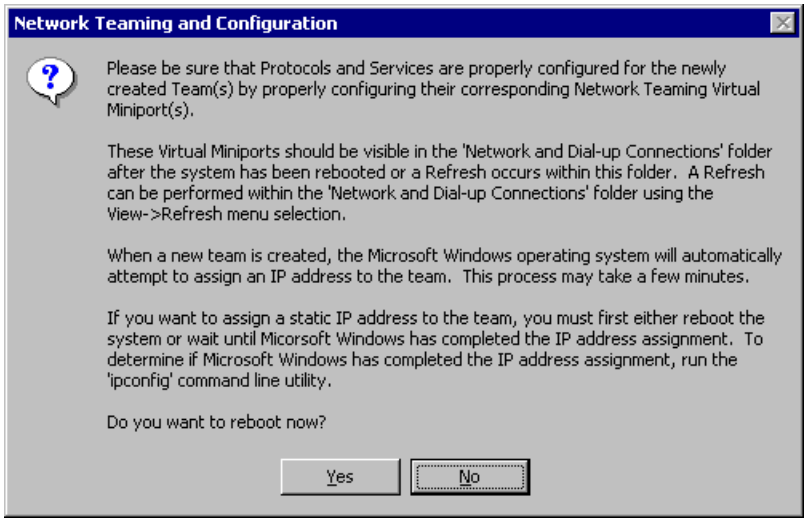

Figure 19: HP Network Teaming dialog box

Configuring the NIC Team Properties

At this point, the NICs are teamed but are not completely configured. Additional procedures include:

- Renaming the teamed connection
- Selecting the option to show an icon on the taskbar
- Configuring TCP/IP on the new team

Renaming the Teamed Connection

The assigned name for the new NIC team connection is “Local Area Connection X,” where X represents the next available connection number generated by the system. HP recommends changing this name to a more meaningful name, such as “NIC Team.”

To change the name of the connection:

1. From the desktop, right-click the **My Network Places** icon, then click **Properties**. The **Network and Dial up Connections** screen is displayed.
2. Move the cursor over each connection icon to view the pop up box of the icon's name. Locate **HP Network Teaming Virtual Miniport**.
3. Right-click the connection icon for **HP Network Teaming Virtual Miniport**, and select **Rename**. Enter a name that is more descriptive than “Local Area Connection X,” such as “NIC Team.”

Showing a Connection Icon on the Taskbar

To show a connection icon:

1. In the **Network and Dial up Connections** screen, double-click the **NIC Team** connection, and then click **Properties**.
2. At the bottom of the screen, select **Show icon in task bar when connected**, and then click **Close**.

Configuring the TCP/IP Protocol on the New Team

After teaming the NICs, a new virtual network adapter for the team is automatically created. However, by default the new adapter is set to DHCP. To manually configure the IP address, perform the following steps.

To enter the TCP/IP address information for the team:

1. From the desktop, go to the **Network and Dial up Connections** screen and click **Properties**. Right-click the **NIC Team** icon and then select **Properties**. A screen similar to the following is displayed.

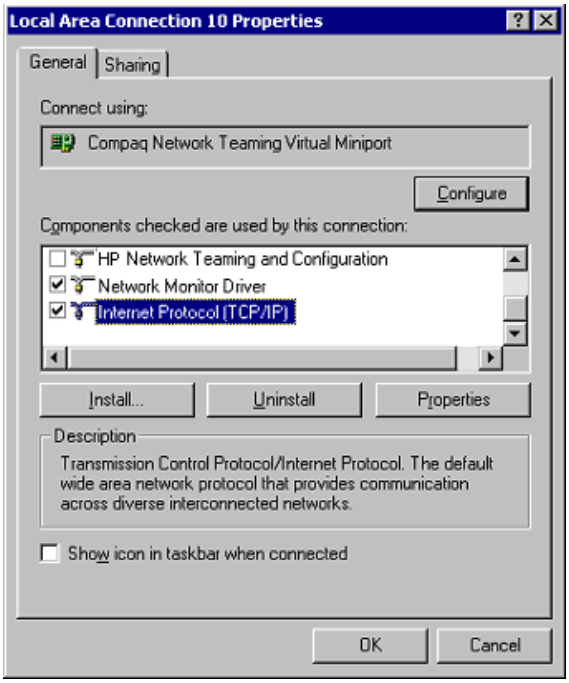

Figure 20: NIC Team Properties dialog box

2. Use the arrows and the scroll bar on the right of the screen to scroll through the **Components** list.
3. Click **Internet Protocol (TCP/IP)** and then click **Properties**. The following screen is displayed:

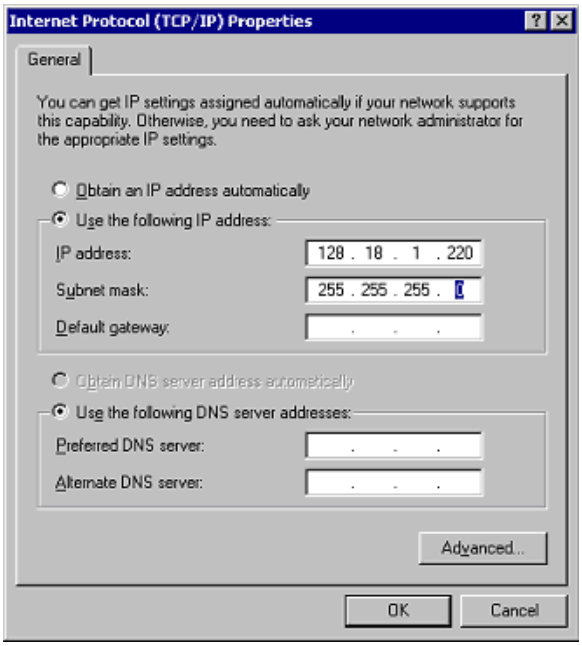

Figure 21: NIC Team TCP/IP Properties dialog box

Note: If a NIC is teamed, do not modify the TCP/IP settings for the individual NIC ports.

4. Select **Use the following IP address**, and enter the IP address and subnet mask. If desired, enter the default gateway.
5. Click **OK**. The Ethernet Team should be working.

Checking the Status of the Team

To check the status of the Ethernet Team, open the HP Network Teaming utility. The **Configuration Properties** screen is displayed, showing the teamed NICs.

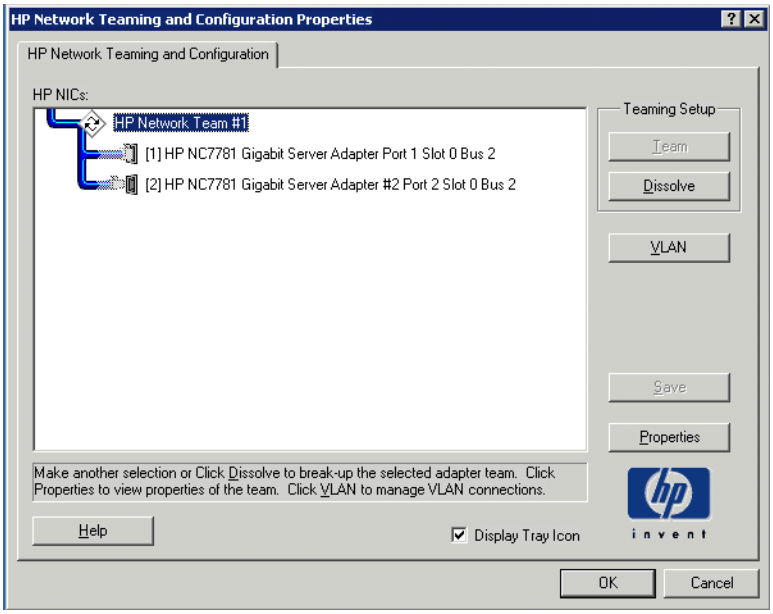

Figure 22: NIC Teaming status

NIC Teaming Troubleshooting

Problems with the NIC teaming feature are diagnosed by the connection icons displayed in the **HP Network Teaming and Configuration** dialog box. The following table lists the error icons for RJ 45 NICs.

Table 2: NIC Teaming Troubleshooting

RJ-45	Description
	Active OK—The NIC is operating properly. The driver is installed in the registry and is loaded. If the NIC is a member of a team, the NIC is active.
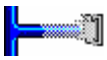	Installed inactive—The NIC is installed and is OK, but is not active.
	Cable fault—The driver is installed in the registry and is loaded. The broken cable indicator means that the cable is unplugged, loose, broken, or the switch or hub is not operating properly. If this icon is displayed, check all network connections and make sure the hub/switch is working properly. When the connection is restored, this icon will change.
	Inactive cable fault—A cable fault has occurred while the NIC was inactive.
	Hardware failure—The driver is installed in the registry and is loaded. The driver is reporting a hardware problem with the NIC. This indicates a serious problem. Contact your HP authorized service provider.
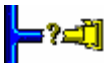	Unknown—The server is unable to communicate with the driver for the installed NIC. The NIC is installed in the registry, but the driver is not. This error occurs when the NIC has been installed but the server has not been restarted. If this problem persists after the server has been restarted, the driver has not been loaded or the Advanced Network Control utility is unable to communicate with the driver. **Note:** Only NICs assigned as members of a team are displayed as Unknown. If a teamed NIC is turned off, it displays as Unknown.
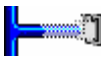	Disabled—The NIC has been disabled through the Device Manager or NCPA.

For more advanced problems with NIC teaming, refer to the help section in the HP Teaming and Configuration utility.

Using Secure Path

Pathing software is required in configurations where multipathing to the storage is desired or required. For clustered products it is highly recommended to maintain two paths to the data as path software allows for datapath failure to occur without forcing a node failover. Secure Path is fully licensed and is contained in the shipping product. Secure Path is installed using the SAN connection tool, found in the HP Utilities tab of the WebUI.

Clustering the NAS Server

Many aspects of configuring a NAS device in a clustered configuration are unique to that environment. The cluster administration chapter later in this guide provides the details behind this specific configuration and the steps necessary to form a cluster. Throughout the remaining chapters, references to the cluster administration chapter are made when special considerations must be applied when utilizing a cluster configuration. Such items include:

- Logical disk support
- Lowest common unit of failover
- File share protocol support
- Users and group management
- Domain considerations
- NFS file share support
- Shadow copies

Managing System Storage

The NAS administrator uses the Array Configuration Utility (ACU) to manage the storage hardware, Disk Manager to manage volumes, and Shadow Copies to manage snapshots. See the following chapters for more detailed information on managing system storage:

- Chapter 3 provides a storage management overview.
- Chapter 4 discusses disk management procedures.
- Chapter 5 discusses snapshot (shadow copy) management procedures.
- Chapter 7 discusses folder and share management procedures.

Creating and Managing Users and Groups

User and group information and permissions determine whether a user can access files. If the NAS device is deployed into a workgroup environment, this user and group information is stored locally on the device. By contrast, if the NAS device is deployed into a domain environment, user and group information is stored on the domain.

To enter local user and group information, see Chapter 6.

Creating and Managing File Shares

Files shares must be set up, granting and controlling file access to users and groups. See Chapter 7 for complete information on managing file shares.

UNIX specific information is discussed in the “Microsoft Services for NFS” chapter.

Storage Management Overview

3

The NAS server is configured at the factory with default system settings and with the NAS operating system installed. No external addressable storage is included with the NAS device. Storage is based on the SAN infrastructure and is configured using the appropriate tools for the particular SAN storage arrays in use.

This chapter defines and discusses physical, logical, and snapshot storage concepts on the HP StorageWorks NAS server.

Additional storage management information is included in the following chapters:

- Chapter 4 discusses disk management procedures.
- Chapter 5 discusses snapshot (shadow copy) management procedures.
- Chapter 7 discusses folder and share management procedures.

Storage Management Process

The lowest level of storage management occurs at the physical drive level. Physical drives are grouped into arrays for better performance and fault tolerance.

The arrays are then configured with RAID fault tolerance and presented to the operating system as logical drives or units, which are called LUNs.

At the virtual level of storage, the WebUI is used to take the LUNs and create basic or dynamic disks, which can then be broken down into partitions or volumes. Folders, subfolders, and file shares are created on the resulting volumes or partitions to organize, store, and give access to the data. The Shadow Copy support is used to create snapshots of the data at specific times.

For organizational and documentation purposes, this administration guide separates physical storage from logical storage.

See Figure 23 for an illustration of these storage management elements.The MSA1000 storage and the associated Array Configuration Utility are used in the example. The Array Storage device in use will impact the devices in use and the relevant storage management software required.

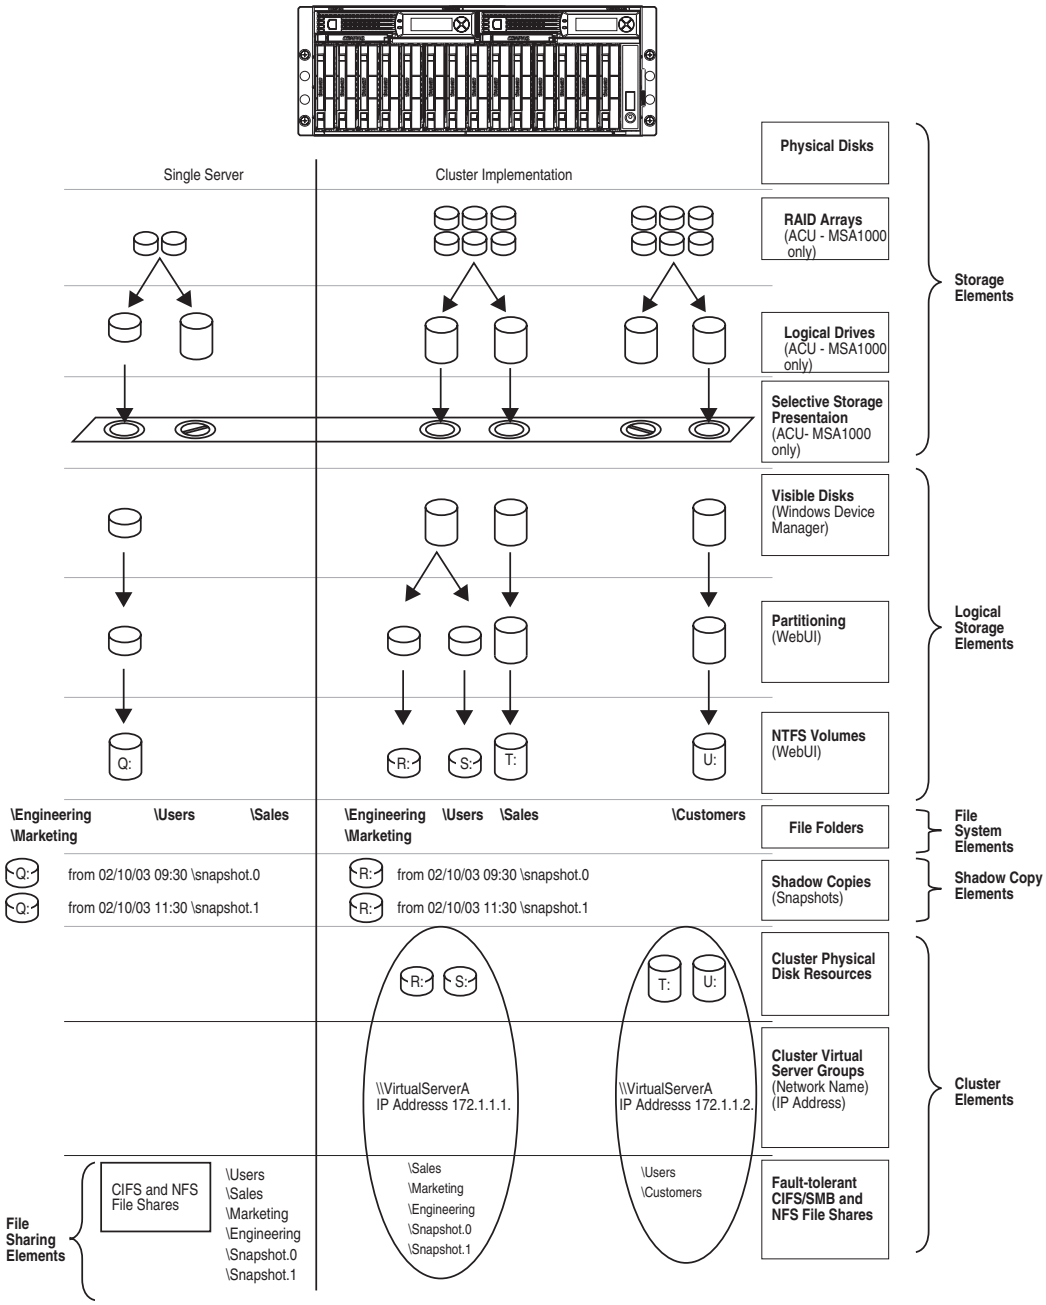

Figure 23: Storage Management process

Storage Elements Overview

The NAS server offers optimized performance for a growing environment. Storage capacity can increase as a business grows without downtime or compromised performance. Storage limitations are based on the type of SAN the NAS server is connected to. See the individual SAN documentation for limitations of Windows Storage Server 2003.

Preliminary physical storage management tasks involve managing:

- Physical Hard Drives
- Arrays
- Logical Drives (LUNs)

Drive array concepts and data protection methods, including fault tolerance options are discussed in this section. This information will help guide decisions on how to best configure the arrays.

Physical Hard Drives

For personal or small business use, the capacity and performance of a single hard drive is adequate. However, larger businesses demand higher storage capacities, higher data transfer rates, and greater security from data loss if drives fail.

Merely adding extra drives to the system increases the total storage capacity, but has little effect on the system efficiency, because data can only be transferred to one hard drive at a time.

Figure 24 illustrates the read/write process with separate physical hard drives.

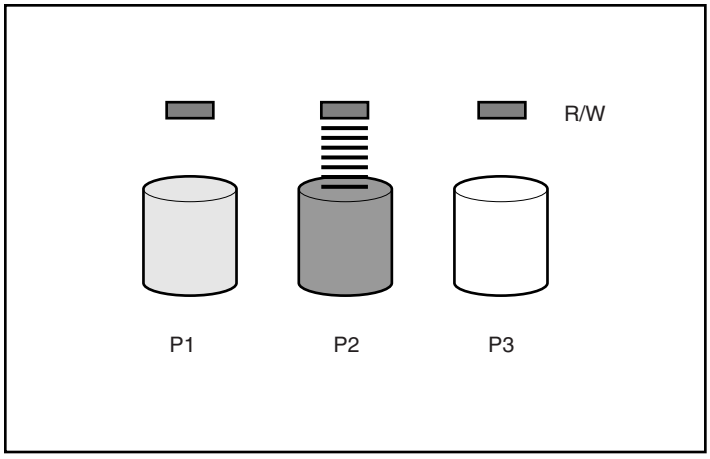

Figure 24: Separate physical drive (P1, P2, P3) read/write (R/W) operations

Arrays

With an array controller installed in the system, the capacity of several physical drives can be logically combined into one or more logical units called arrays. When this is done, the read/write heads of all the constituent physical drives are active simultaneously, dramatically reducing the overall time required for data transfer.

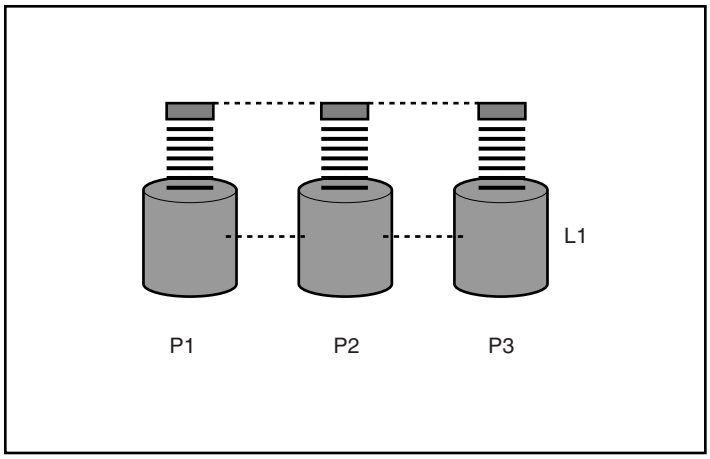

Figure 25: Configuring the physical drives into an array dramatically improves read/write efficiency

Because the read/write heads are active simultaneously, the same amount of data is written to each drive during any given time interval. Each unit of data is termed a block. The blocks form a set of data stripes over all the hard drives in an array, as shown in Figure 26.

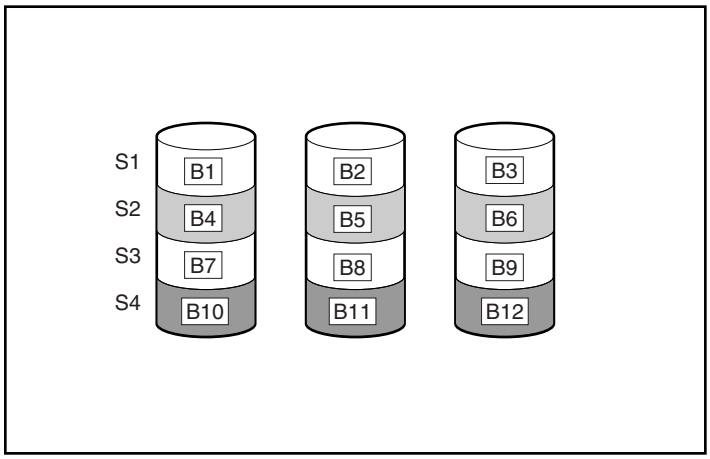

Figure 26: RAID 0 (data striping) (S1-S4) of data blocks (B1-B12)

For data in the array to be readable, the data block sequence within each stripe must be the same. This sequencing process is performed by the array controller, which sends the data blocks to the drive write heads in the correct order.

A natural consequence of the striping process is that each hard drive in a given array will contain the same number of data blocks.

Note: If one hard drive has a larger capacity than other hard drives in the same array, the extra capacity is wasted because it cannot be used by the array.

Logical Drives (LUNs)

As previously stated, drive array technology distributes data across a series of individual hard drives to unite these physical drives into one or more higher performance arrays. Distributing the data allows for concurrent access from multiple drives in the array, yielding faster I/O rates than non arrayed drives.

While an array is a physical grouping of hard drives, a logical drive is the configuration of the arrays that is presented to the operating system.

When planning to allocate space on the NAS device, consider that the maximum number of LUNs in a dynamic disk is 32 and the largest single LUN that can be utilized by the operating system is 2 TB. It should also be noted that the largest basic disk that can exist is 2 TB and the largest volume that can exist is 64 TB. Format of the partition or volume impacts the largest file system that can exist as well.

After the physical drives are grouped into arrays, they are ready to be converted into logical drives. Options for working with arrays vary from SAN storage to SAN storage system. The individual documentation included with each storage system should be reviewed.

It is important to note that a LUN may extend over (span) all physical drives within a storage controller subsystem, but cannot span multiple storage controller subsystems.

After a LUN has been created, it is possible to extend the size. See the section “Dynamic Growth” in chapter 4 for additional information on LUN extension and use by the operating system.

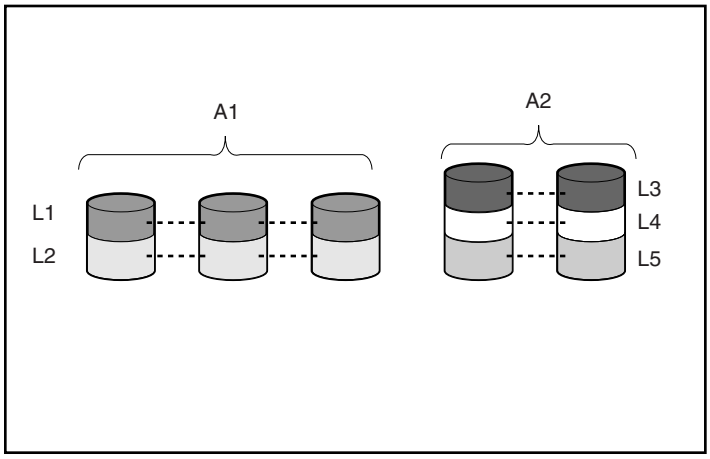

Figure 27: 2 arrays (A1, A2) and 5 logical drives (L1 through L5) spread over 5 physical drives

Note: This type of configuration may not apply to all supported SANs and serves only as an example.

Drive failure, although rare, is potentially catastrophic. For example, in the previous figure using simple striping, failure of any hard drive will lead to failure of all logical drives in the same array, and hence to data loss.

To protect against data loss from hard drive failure, arrays should be configured with fault tolerance. Several fault tolerance methods have been devised and are described in the following sections.

Fault-Tolerance Methods

Different RAID (redundant array of independent disks) types use different methods of striping the arrays and different ways of writing data and parity to the drives to offer a variety of fault tolerance and capacity usage. The RAID methods supported by the NAS server include:

- RAID 0—Data Striping only, no fault tolerance
- RAID 1+0—Drive Mirroring and striping
- RAID 5—Distributed Data Guarding
- RAID ADG—Advanced Data Guarding (ADG)

Further protection against data loss can be achieved by assigning an online spare to an array. This hard drive contains no data and is contained within the same storage subsystem as the other drives in the array. When a hard drive in the array fails, the controller can then automatically rebuild information that was originally on the failed drive onto the online spare. This quickly restores the system to full RAID level fault tolerance protection.

Note: The ADG feature is available only with the MSA1000. RAID 5DP is available only with HP Virtual Arrays and is equivalent to ADG.

These fault tolerance methods are discussed in the following paragraphs.

RAID 0—Data Striping

This configuration provides striping of the array to improve read and write performance, but offers no redundancy of data and therefore no protection against data loss when a drive fails. However, RAID 0 is useful for rapid storage of large amounts of noncritical data (for printing or image editing, for example) or when cost is the most important consideration.

When creating RAID 0 arrays, carefully consider how many drives to include in the array. Statistically, the chance of a drive failure increases with each additional drive that is included in an array. Based upon laboratory testing, HP recommends including no more than 7 drives in a RAID 0 array.

See Figure 26 for an illustration of the data striping technique.

Advantages

- Highest performance method for reads and writes
- Lowest cost per unit of data stored
- All drive capacity is used to store data; none is used for fault tolerance

Disadvantages

- All data on logical drive is lost if a hard drive fails
- Cannot use an online spare
- Data can only be preserved by being backed up to external media

RAID 1+0—Drive Mirroring and Striping

In this configuration, information on one drive is duplicated onto a second drive, creating identical copies of the information as shown in Figure 28. Therefore, this method provides the best fault tolerance. RAID 1+0 requires an even number of drives and is the only method for fault tolerance protection if only two drives are installed or selected for an array. If more than two drives are in an array, the data is striped across all of the drives in the array.

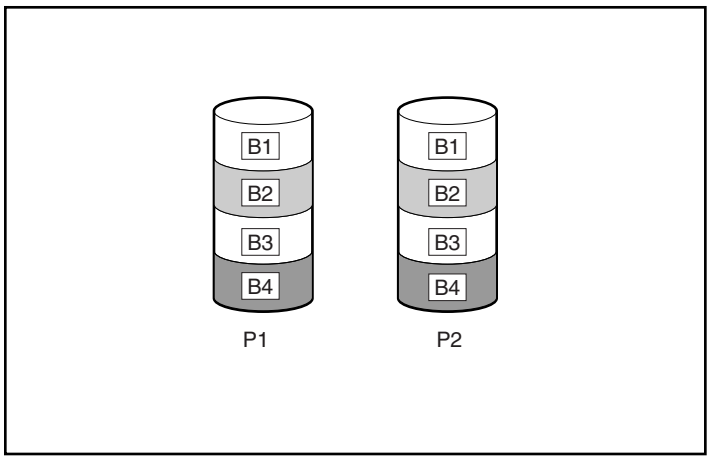

Figure 28: RAID 1+0 (drive mirroring) of P1 onto P2

This method is useful when high performance and data protection are more important than the cost of hard drives. The operating system drives are mirrored. If one drive fails, the mirror drive immediately takes over and normal system operations are not interrupted.

Note: HP supports a configuration that uses RAID 1+0 on the system drives in a two drive RAID array.

Caution: If two drives being mirrored to each other both fail, data loss occurs.

Advantages

Drive mirroring offers:

- The highest read and write performance of any fault-tolerant configuration.
- Protection against data loss if one drive fails.
- Data preservation in a RAID 1+0 system, when more than one drive fails, as long as none of the failed drives are mirrored to another failed drive.

Disadvantages

Some disadvantages of drive mirroring are:

- Increased expense—Since many drives must be used for fault tolerance and hard drives must be added in pairs.
- Decreased storage capacity—Only 50% of the total drive capacity is usable.

RAID 5—Distributed Data Guarding

Using this method, a block of parity data (rather than redundant data) is calculated for each stripe from the data that is in all other blocks within that stripe. The blocks of parity data are distributed over every hard drive within the array, as shown in the figure below. When a hard drive fails, data on the failed drive can be rebuilt from the parity data and the user data on the remaining drives. This rebuilt data can be written to an online spare.

This configuration is useful when cost, performance, and data availability are equally important.

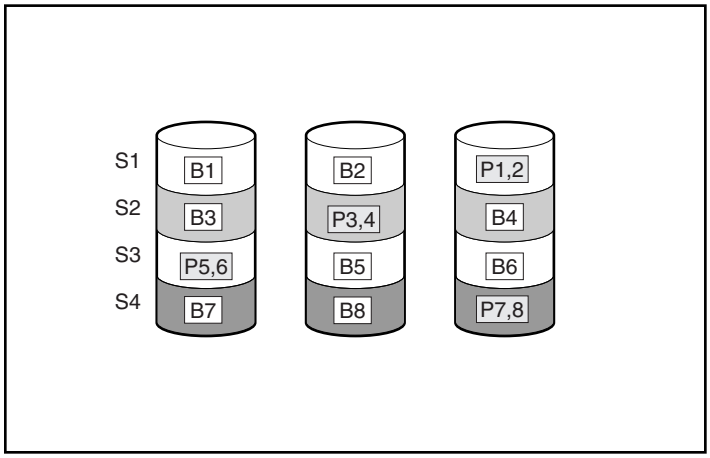

Figure 29: RAID 5 (distributed data guarding) showing parity information (P)

Spreading the parity across all the drives allows more simultaneous read operations and higher performance than data guarding (RAID 4). If one drive fails, the controller uses the parity data and the data on the remaining drives to reconstruct data from the failed drive. RAID 5 allows the system to continue operating with reduced performance until the failed drive is replaced. However, if more than one drive fails, RAID 5 also fails and all data in the array is lost.

Distributed data guarding uses the equivalent of one drive to store parity information and requires an array with a minimum of three physical drives. In an array containing three physical drives, distributed data guarding uses 33 percent of the total logical drive storage capacity for fault tolerance; a 14 drive configuration uses seven percent.

Note: Given the reliability of a particular generation of hard drive technology, the probability of an array experiencing a drive failure increases with the number of drives in an array. HP recommends the number of drives in a RAID 5 array not exceed 14.

Advantages

Distributed data guarding offers:

- High read and write performance.
- Protection against data loss if one drive fails.
- Increased usable storage capacity, since capacity equal to only one physical drive is used to store parity information.

Disadvantages

Some disadvantages of distributed data guarding are:

- Lower write performance than RAID 0 or RAID 1+0.
- Possibility of data loss if a second drive fails before data from the first failed drive has been rebuilt.

RAID ADG—Advanced Data Guarding and RAID 5DP—Double Parity

RAID ADG and RAID 5DP are similar to RAID 5 in that parity information is generated (and stored) to protect against data loss caused by drive failure. With RAID ADG and RAID 5DP, however, two different sets of parity data are used. This allows data to still be preserved if two drives fail. As can be seen from Figure 30, each set of parity data uses up a capacity equivalent to that of one of the constituent drives, for a total parity usage of two drives of space.

This method is most useful when data loss is unacceptable, but cost must also be minimized. The probability that data loss will occur when configured with RAID ADG or RAID 5DP is less than when configured with RAID 5.

Note: The ADG feature is available only with the MSA1000. RAID 5DP is available only with HP Virtual Arrays and is equivalent to ADG.

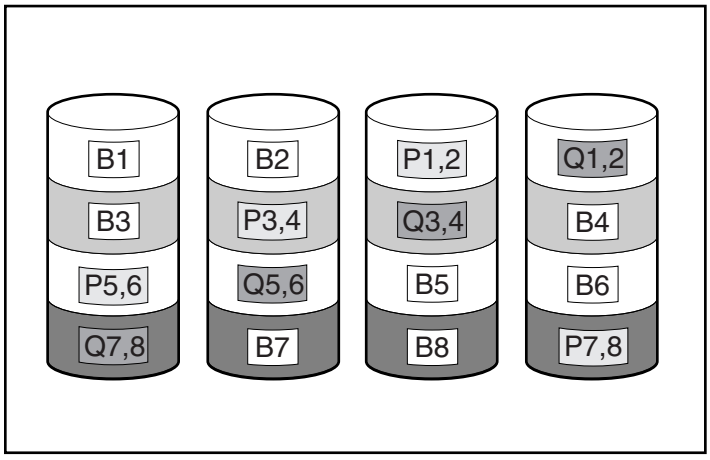

Figure 30: RAID ADG (advanced data guarding) with two sets of parity data

Advanced Data Guarding technology offers the best combination of fault tolerance and usable disk space among RAID levels.

This technology allows the safe deployment of large capacity disk drives and the creation of very large storage volumes without expensive overhead to protect business critical data. This technology provides more flexibility in responding to drive failures without the fear of costly server downtime.

Advance Data Guarding protects against multiple disk failures, while requiring the capacity of two drives in an array of up to 56 disk drives to be set aside for dual sets of distributed parity data. It provides data protection greater than RAID 5, and also has the capacity utilization efficiency similar to RAID 5.

Advantages

- High read performance.
- High data availability-any two drives can fail without loss of critical data.

Disadvantages

- Relatively low write performance (lower than RAID 5), due to the need for two sets of parity data.

The table below summarizes the important features of the different kinds of RAID supported by the Smart Array controllers. The decision chart in the following table may help determine which option is best for different situations.

Table 3: Summary of RAID Methods

	RAID 0 Striping (no fault tolerance)	RAID 1+0 Mirroring	RAID 5 Distributed Data Guarding	RAID ADG Advanced Data Guarding
Maximum number of hard drives	N/A	N/A	14	Storage system dependent
Tolerant of single hard drive failure?	No	Yes	Yes	Yes
Tolerant of multiple simultaneous hard drive failure?	No	For RAID 1+0, if the failed drives are not mirrored to each other	No	Yes (two drives can fail)

Online Spares

Further protection against data loss can be achieved by assigning an online spare (or hot spare) to any configuration except RAID 0. This hard drive contains no data and is contained within the same storage sub system as the other drives in the array. When a hard drive in the array fails, the controller can then automatically rebuild information that was originally on the failed drive onto the online spare. This quickly restores the system to full RAID level fault tolerance protection. However, unless RAID ADG is being used, which can support two drive failures in an array, in the unlikely event that a third drive in the array should fail while data is being rewritten to the spare, the logical drive will still fail.

Physical Storage Best Practices

Minimally, choosing the best disk carving strategy includes the following policies:

- Analyze current corporate and departmental structure.
- Analyze the current file server structure and environment.
- Plan properly to ensure the best configuration and use of storage.
  - Determine the desired priority of fault tolerance, performance, and storage capacity.
  - Use the determined priority of system characteristics to determine the optimal striping policy and RAID level.
- Include the appropriate number of physical drives in the arrays to create LUNs of desired sizes.

Logical Storage Elements Overview

Logical Storage elements consist of those components that translate the physical storage elements to the file system elements as presented in Figure 23. The NAS server utilizes the WebUI to manage the various types of disk presented to the file system. The WebUI has two types of LUN presentation, basic disk and dynamic disk. Each of these types of disk has special features that enable different types of management. Through the use of basic disks, primary partitions or extended partitions may be created. Partitions can only encompass one LUN. Through the use of dynamic disks, volumes can be created that span multiple LUNS. The WebUI can be used to convert disks to dynamic and back to basic, and manage the volumes residing on dynamic disks. Other options include the ability to delete, extend, mirror, and repair these elements.

The sections below briefly discuss each of these types of representations and the considerations that need to be observed.

More detailed information regarding the WebUI for disk management activities can be obtained in the Disk Management Chapter.

Partitions

Partitions exist as either Primary Partitions or Extended Partitions and can be composed of only one Basic disk no larger than 2 TB. Basic disks can also only contain up to four primary partitions, or three primary partitions and one extended partition. In addition, the partitions on them cannot be extended beyond the limits of a single LUN. Extended partitions allow the user to create multiple logical drives. These partitions or logical disks can be assigned drive letters

or be mounted as mount points on existing disks. If mount points are utilized, it should be noted that Services for UNIX does not support mount points at this time. The use of mount points in conjunction with NFS shares is not supported.

Volumes

When planning dynamic disks and volumes there is a limit to the amount of growth a single volume can undergo. Volumes are limited in size and are limited to no more than 32 separate LUNs with each LUN not exceeding 2 terabytes (TB). Volumes also cannot exceed 64 TB of disk space.

The RAID level of the LUNs included in a volume must be considered. All of the units that make up a volume should have the same high-availability characteristics. In other words, the units should all be of the same RAID level. For example, it would be a bad practice to include both a RAID 1+0 and a RAID 5 array in the same volume set. By keeping all the units the same, the entire volume retains the same performance and high-availability characteristics, making managing and maintaining the volume much easier. It should be noted that if a dynamic disk goes offline, then the entire volume dependent on the one or more dynamic disks is unavailable. There could be a potential for data loss depending on the nature of the failed LUN.

Volumes are created out of the dynamic disks and can be expanded on the fly to extend over multiple dynamic disks if they are spanned volumes. However, once a type of volume is selected it cannot be altered. For example, a spanning volume cannot be altered to a mirrored volume without deleting and recreating the volume, unless it is a simple volume. Simple volumes can be mirrored or converted to spanned volumes. Fault tolerant disks cannot be extended either. Therefore, selection of the volume type is important. Please note that the same performance characteristics on numbers of reads and writes apply when using fault tolerant configurations as is the case with controller based RAID. These volumes can also be assigned drive letters or be mounted as mount points off existing drive letters. In general, HP recommends utilizing the Array controller for the management of fault tolerance over the use of Windows Storage Server 2003 software RAID since it places an additional level of operating system overhead on volumes. If mount points are utilized, it should be noted that Services for UNIX does not support mount points at this time.

The administrator should carefully consider how the volumes will be carved up and what groups or applications will be using them. For example, putting several storage-intensive applications or groups into the same dynamic disk set would not be efficient. These applications or groups would be better served by being divided up into separate dynamic disks, which could then grow as their space requirements increased, within the allowable growth limits.

Note: Dynamic disks cannot be used for clustering configurations because Microsoft Cluster only supports basic disks.

Utilizing Storage Elements

No matter which type of storage element is created in the WebUI the last step in creating the element is determining its drive letter or mount point and formatting the element. Each element created can exist as a drive letter(s), assuming one is available and/or as mount points off of an existing folder of a drive letter. Either method is supported. However, mount points

can not be utilized for shares that will be shared using Microsoft Services for Unix (NFS). They can be setup with both but the use of the mount point in conjunction with NFS shares causes instability with the NFS shares.

Formats consist of NTFS, FAT32, and FAT and all three types can be used on the NAS device. However, the Volume Shadow Copy Service can only utilize volumes that are NTFS formatted. Also, quota management is possible only on NTFS.

Volume Shadow Copy Service Overview

The Volume Shadow Copy Service (VSS) provides an infrastructure for creating point-in-time snapshots (shadow copies) of volumes. VSS supports 64 shadow copies per volume.

Shadow Copies of Shared Folders resides within this infrastructure and helps alleviate data loss by creating shadow copies of files or folders that are stored on network file shares at pre-determined time intervals. In essence, a shadow copy is a previous version of the file or folder at a specific point in time.

By using shadow copies, a NAS server can maintain a set of previous versions of all files on the selected volumes. End users access the file or folder by using a separate client add-on program, which enables them to view the file in Windows Explorer.

Shadow copies should not replace the current backup, archive, or business recovery system, but they can help to simplify restore procedures. For example, shadow copies cannot protect against data loss due to media failures, however, recovering data from shadow copies can reduce the number of times needed to restore data from tape.

Detailed information on Shadow Copies can be found in Chapter 5 of this guide.

File System Elements

File system elements are composed of the folders and subfolders that are created under each Logical Storage element (partitions, logical disks, and volumes). Folders are used to further subdivide the available file system providing another level of granularity for management of the information space. Each of these folders can contain separate permissions and share names that can be used for network access. Folders can be created for individual users, groups, projects, and so on.

Detailed information on file system elements can be found in Chapter 7 of this guide.

File-Sharing Elements

The NAS server supports several file sharing protocols, including DFS, NFS, FTP, HTTP, and Microsoft SMB. On each folder or Logical Storage element, different file sharing protocols can be enabled using specific network names for access across a network to a variety of clients. Permissions can then be granted to those shares based on users or groups of users in each of the file sharing protocols.

Detailed information on file-sharing elements can be found in Chapter 7 of this guide.

Clustered Server Elements

The HP StorageWorks NAS server supports several file-sharing protocols, including DFS, NFS, FTP, HTTP, and Microsoft SMB. Only NFS, FTP, and Microsoft SMB are cluster-aware protocols. HTTP can be installed on each node but the protocols cannot be set up through cluster administrator, nor will they failover during a node failure.

Caution: AppleTalk shares should not be created on clustered resources as this is not supported by Microsoft Clustering and data loss may occur.

Network names and IP address resources for the clustered file share resource may also be established for access across a network to a variety of clients. Permissions can then be granted to those shares based on users or groups of users in each of the file sharing protocols.

Disk Management

4

Disk Management is core to the Windows NAS product. The process of creating storage elements and presenting them to the NAS OS is facilitated by the use of the WebUI. This chapter documents the contents of the WebUI for disks and volume management.

WebUI Disks Tab

The online Storage Guide provides an overview of the storage management process as a supplement to this administration guide.

The primary web page for facilitating disks and volume creation is illustrated in Figure 31. The figure in the diagram illustrates the process of creating arrays, volumes, and shadow copies. The diagram on the left illustrates the logical steps used to manage disks, beginning with the array management at the top. The process follows the diagram from top to bottom and the selectable menu items from left to right on the page:

1. Create arrays and LUNS via the appropriate storage array management software
2. Create disks via the WebUI
3. Create volumes via the WebUI

To manage disks and volumes via the WebUI, click on **Disks**.

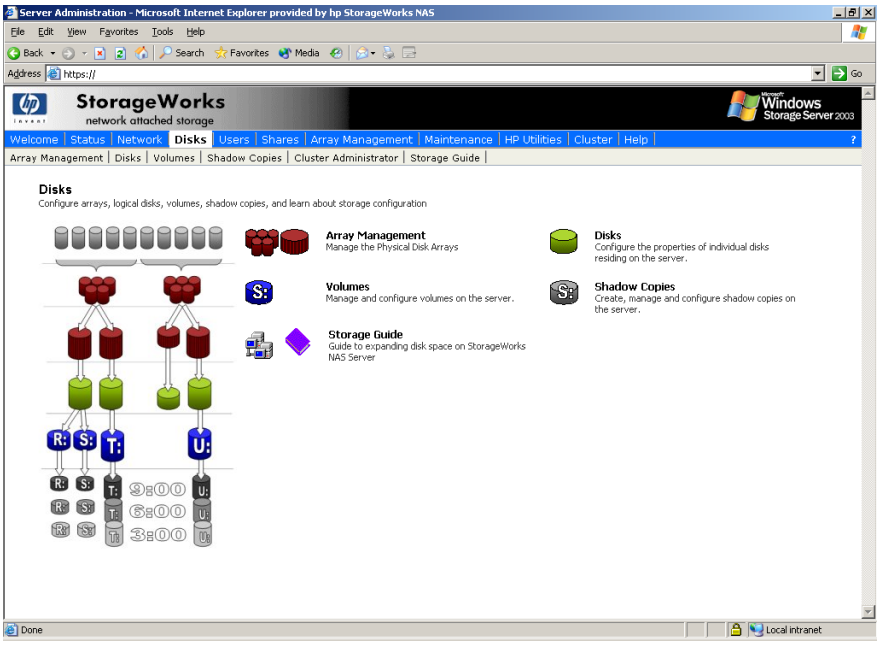

Figure 31: Disks menu

The Disks tab contains the following task items for configuring the NAS device:

Table 4: Disks Tab Options

Option	Task
Array Management	Open the Array Management screen to access the ACU and links to other storage array management elements.
Disks	Manage logical disks. Observe disk capacity and status, scan for new disks, view detailed disk properties, and create new volumes.
Volumes	Manage disk space usage by enabling quotas, scheduling disk defragmentation, and performing detailed volume management using the Manage button.
Shadow Copies	Manage shadow copies of shared folders on the volume. Shadow copies are read-only copies of shared data that provide users via network shares with a way to view, and, if necessary, restore previous versions or deleted files.
Storage Guide	Provides a detailed list of the procedures required to configure and create disks and volumes on NAS devices.

Storage Configuration Overview

Physical disks can be arranged as RAID arrays for fault tolerance and enhanced performance, then segmented into logical disks of appropriate sizes for particular storage needs. These logical disks then become the volumes that appear as drives on the NAS server.

Note: This type of configuration may not apply to all supported SANs and serves only as an example.

The following steps are an example of a storage configuration using an HP StorageWorks MSA1000.

Step 1: Create Disk Arrays

1. Click **Array Management** on the **Disks** tab.
2. Click **Array Configuration Utility** and log in to the management page in another browser window.

   The Array Configuration Utility will start.
3. Select the proper array controller in the left pane of the interface before beginning array configuration. Some NAS systems are equipped with array controllers for both internal and external storage.

Consult the Help available in ACU for details on creating arrays, if necessary.

Step 2: Create Logical Disks from the Array Space

From the ACU:

1. Select a previously created array.
2. Click **Create Logical Drive** from the right pane of the ACU.
3. Complete the logical drive creation wizard to designate some or all of the array space as a logical disk.

Depending on how many physical disks are included in the array, several different types of logical disks are possible. Consult the ACU Help for details on creating logical disk drives.

Step 3: Verify newly created logical disks

1. Click **Disks** on the **Disk** tab.
2. Verify that disks matching the newly created sizes are displayed.
3. Click on **Initialize Disk** to initialize the Disk.

Note: By default the disk is basic. Click Convert Disk to make the disk dynamic.

Step 4: Create a Volume on the new logical disk

1. Click **Create New Volume**.
2. Enter the volume size.
3. Select a drive letter.
4. Enter a mount point, if desired.
5. Select to format the volume, if desired.
6. Click **OK**.
7. Select whether or not to quick format the volume.
8. Enter a volume label.
9. Enter the allocation unit size.
10. Click **OK**.

Array Configuration Utility (MSA1000 and internal OS drives only)

RAID arrays and LUNs are created and can be managed using the HP Array Configuration Utility.

Features of ACU:

- Graphical representation of drive array configurations with wizards that help optimize array configuration
- Online spare (hot spare) configuration
- Separate fault tolerance configurations on a logical drive (LUN) basis
- Easy capacity expansion of arrays
- Online RAID level and stripe size migration
- Manages OS and data drives

Each time the Array Configuration Utility is run, it analyzes the configuration of the Array Controllers installed in the system. From the Main page various options are available to change or reconfigure the controller(s). This document only covers a subset of the functions available in the ACU. For complete documentation on ACU, refer to the comprehensive online help found within the ACU tool.

Using the ACU to Configure Storage

To configure storage:

1. Open the WebUI and navigate to the **Disks** tab.
2. Click **Array Management**.

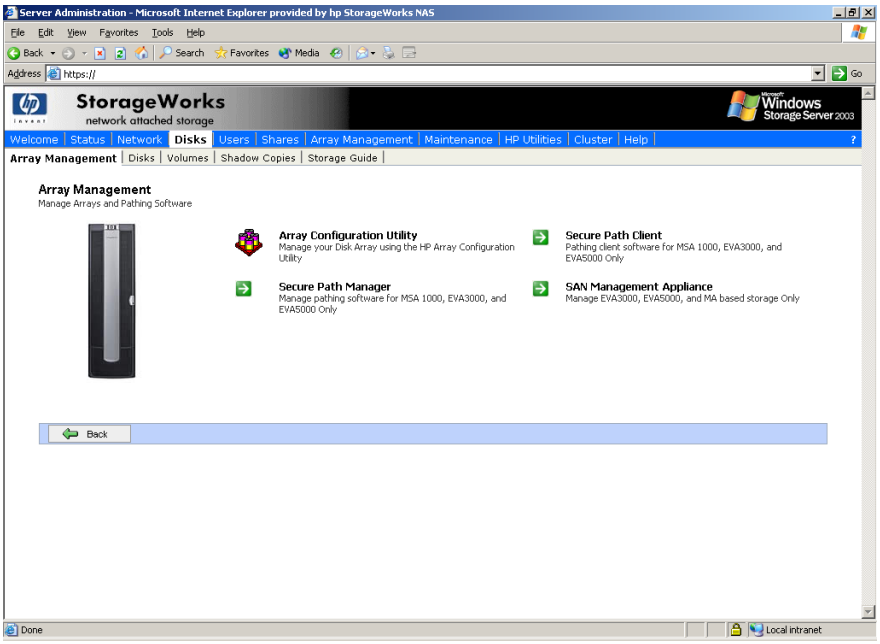

Figure 32: Array Management screen

3. Click **Array Configuration Utility**.

Note: ACU is used to manage and configure array-based storage.

4. Log in to the ACU utility. The default user name is `administrator` and the default password is `administrator`.

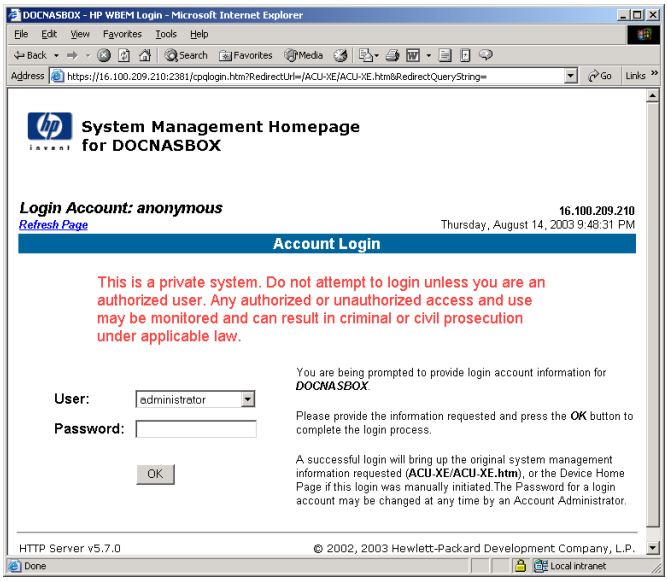

Figure 33: Systems Management Homepage

The Array Configuration Utility is displayed.

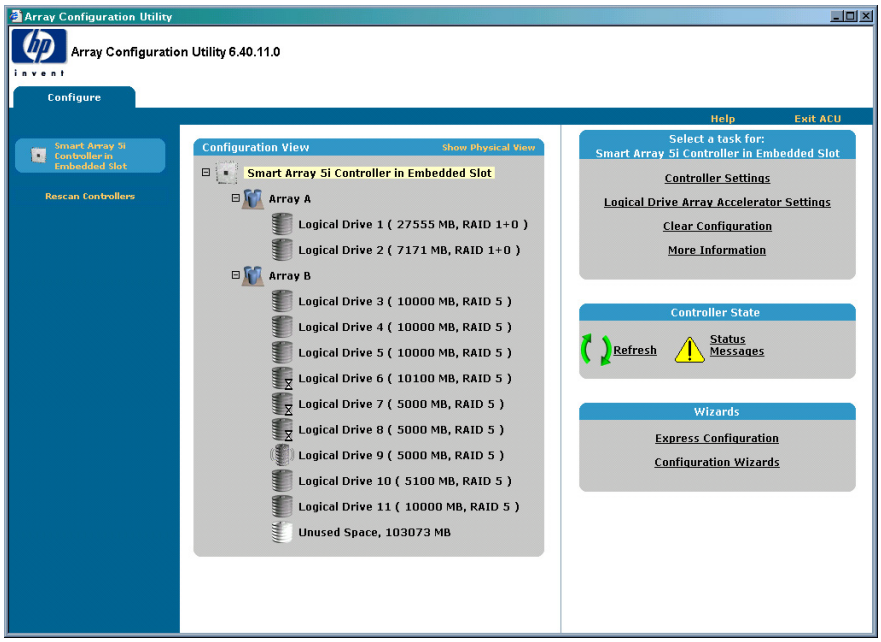

Figure 34: Array Configuration Utility

5. Select a controller from the list on the left side to begin configuration.

- The controller named Smart Array 5i Controller in the embedded slot is for all drives in the server chassis, and drives contained in an external storage enclosure attached to the Smart Array 5i on the server head if present.
- Additional controllers (if present) are used for all externally SCSI attached storage.

Caution: On the Smart Array 5i controller there are two logical drives pre-configured under Array A. These two logical drives are configured for the NAS operating system and should not be altered.

6. After the controller is selected there are three ways to configure the storage:
   - **Express Configuration**

     Select **Express Configuration** to be asked a few simple questions and allow the controller to be configured automatically. The **Express Configuration** is the easiest and fastest way to configure a controller and will provide the most reasonable configuration possible.
   - **Configuration Wizards**

     Select **Configuration Wizards** to configure a controller through a series of wizards, which provides a guide through the configuration process. Choosing **Configuration Wizards** is not the fastest or easiest way to configure a controller, but it does offer more control over the configuration and provides for a more individualistic setup.
   - **Standard Configuration (default)**

     Select **Standard Configuration** to quickly configure a controller. Choosing **Standard Configuration** is the fastest way to configure a controller but requires an intermediate to advanced level of knowledge concerning storage. The Standard Configuration path offers the least amount of help or step-by-step guides and does not provide a FAQ panel, assuming the user knows exactly what they would like to accomplish and are very familiar with the concepts required to complete the task.

     The default method is the standard configuration method. The steps that follow are for creating an array using the standard configuration mode.
7. Click **Create Array**.
8. Select all of the drives to be included in the array and click **OK**.
9. Select the array that was just created and click on **Create logical Drive** at the right.
10. Select the desired Fault Tolerance, Stripe Size, and Size of the logical disk, and click **OK**.

    The Fault tolerance level depends on the amount of disks selected when the array was created. A minimum of two disks is required for a RAID 0+1 configuration, three disks for a RAID 5 configuration, and four disks for a RAID 5 ADG configuration.
11. After all logical disks have been created, click **Save**.
12. Click **Exit ACU** to exit the ACU session.

ACU Guidelines

- Do not modify Array A off of the Smart Array 5i controller as it contains the NAS OS
- Spanning more than 14 disks with a RAID 5 volume is not recommended
- Designate spares for RAID sets to provide greater protection against failures
- RAID sets cannot span controllers
- A single array can contain multiple logical drives of varying RAID settings.
- Extending and Expanding Arrays and Logical Drives is supported
- RAID migration is supported

Managing Disks

From the **Disks** tab of the WebUI, select **Disks**. The page displays the physical disks that are associated with the NAS device and the volumes that are created on them. Multiple volumes may appear on multiple disks depending on whether the volumes are simple, spanned, or multi-volumes/partitions exist. The page also displays the type of disk (basic or dynamic).

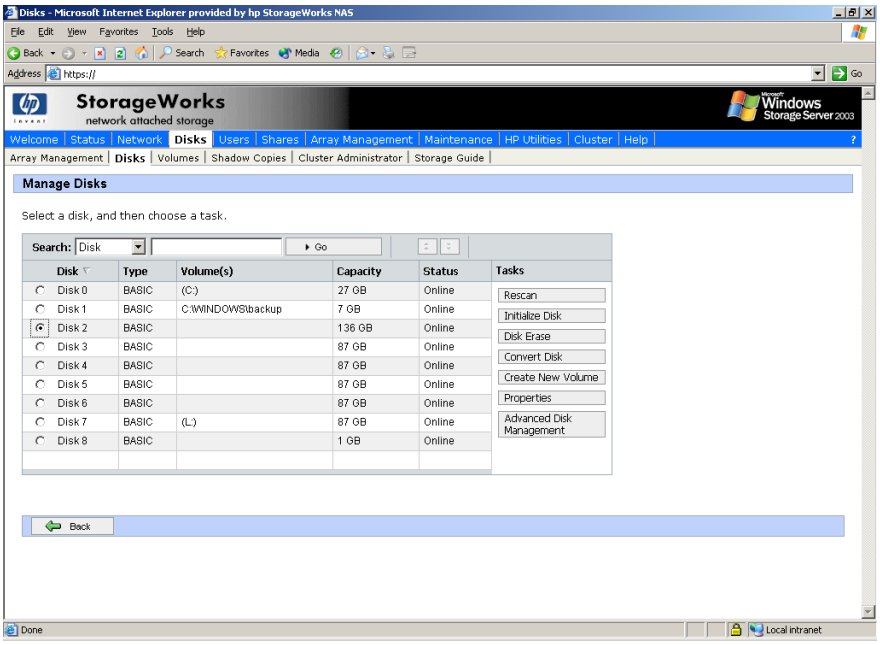

Figure 35: Manage Disks screen

Table 5: Manage Disks Options

Option	Task
Rescan	Select to detect a new disk added to the system. By default, drives are dynamically recognized by the system. Occasionally a rescan of the hardware is required. The rescan is not synchronous and may require a browser refresh after the scan is initiated to display the new content.
Initialize Disk*	Initializes any empty disk to type basic.
Disk Erase*	Erases the selected disk.
Convert Disk*	Converts the selected disk from basic to dynamic, or dynamic to basic.
Create New Volume	Select to create a new volume.
Properties	Select to display the properties of the selected disk.
Advanced Disk Management	Select to open the Disk Management utility and perform advanced disk management tasks. Please see the online Disk Management help pages for complete documentation.
* These tasks cannot be completed on clustered resources.	

Creating a New Volume via the WebUI

To create a new volume via the WebUI:

1. Click the **Disks** tab, then click **Disks**.
2. Select the Disk to create the new volume on.
3. Click **Create New Volume**.

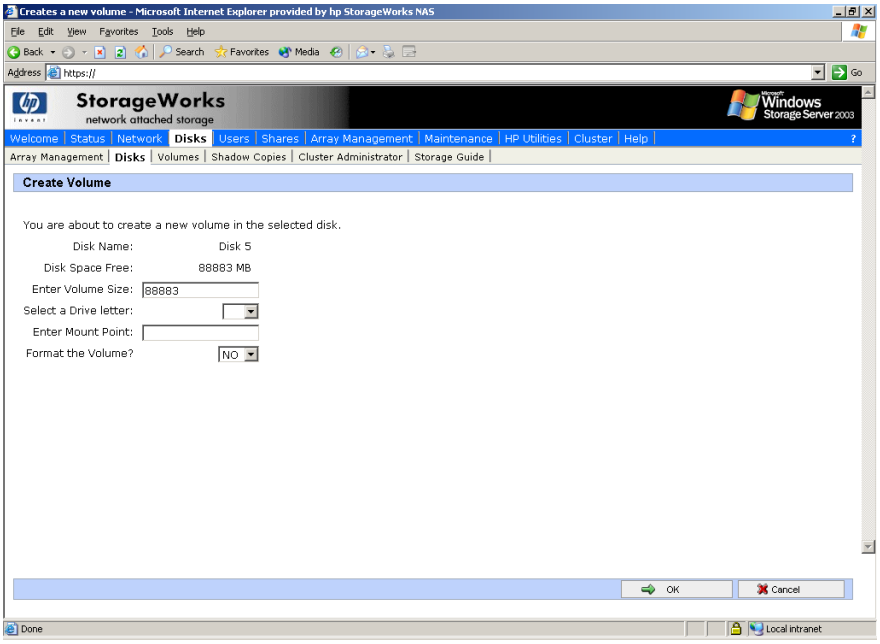

Figure 36: Creating a new volume, page 1

4. Enter the volume size.
5. Select a drive letter.
6. Enter a mount point, if desired.
7. Select to format the volume, if desired.
8. Click **OK**.

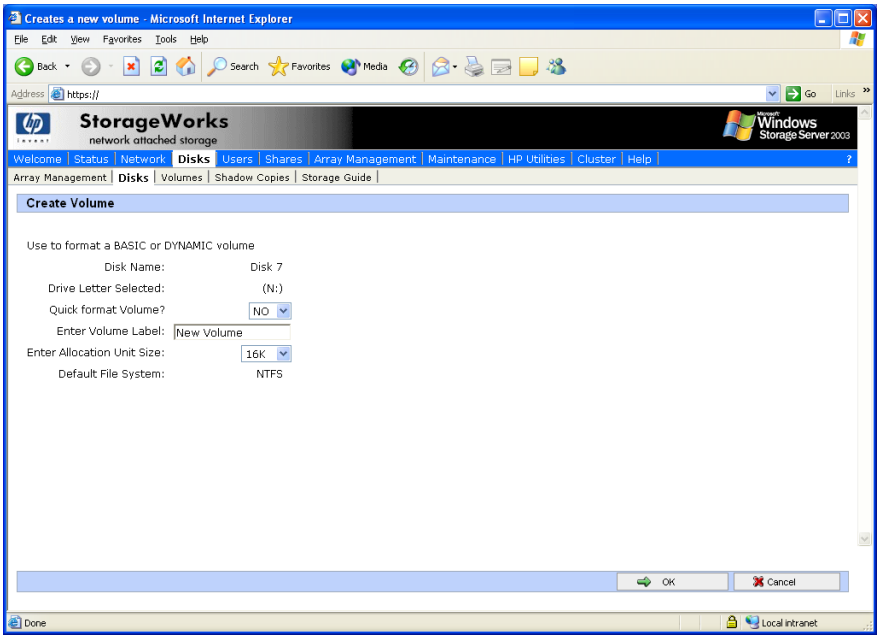

Figure 37: Creating a new volume, page 2

9. Select whether or not to quick format the volume.
10. Enter a volume label.
11. Enter the allocation unit size.
12. Click **OK**. The Manage Disks page is displayed.

Advanced Disk Management

When the Advanced Disk Management button on the Manage Disks screen is selected, the Disk Management Utility is opened in a remote desktop session. The Disk Management tool is a system utility for managing hard disks and the volumes, or partitions, that they contain. The WebUI provides most of the functionality required for NAS disk management. However there are some instances where the Disk Manager is required. For example, to reassign a drive letter or mount point or to create software based RAID fault-tolerant disk systems.

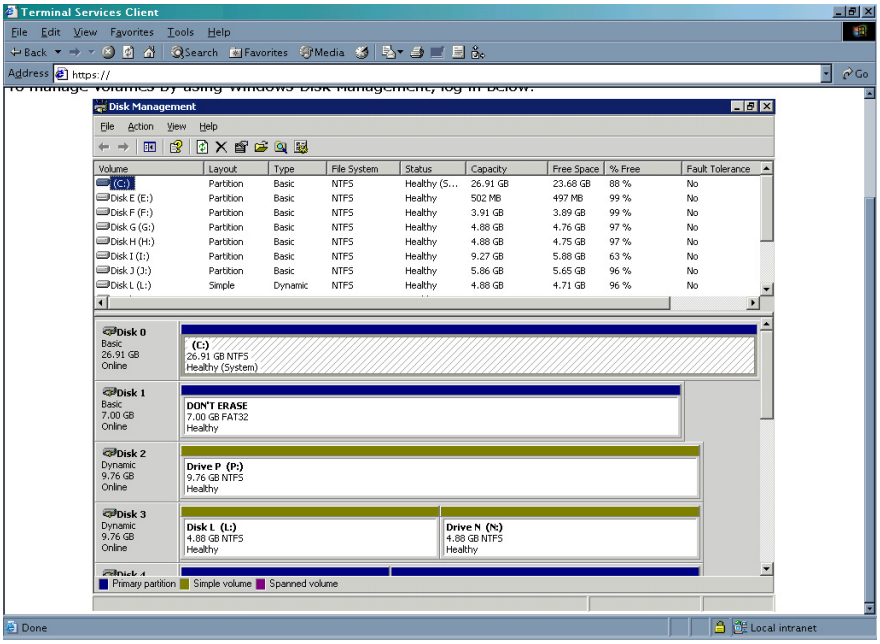

Figure 38: Disk Management utility

Note: When the Disk Management utility is accessed, the Remote Desktop connection assumes a dedicated mode and can only be used to manage disks and volumes on the server. Navigating to another page during an open session closes the session.

Note: It may take a few moments for the Remote Desktop Connection session to log off when closing Disk Management.

Guidelines for Managing Disks

When managing disks and volumes:

- Read the online Help found in the WebUI.
- Do not alter the Operating System Disk Labeled Local Disk C:.
- Do not alter the disk labeled “DON’T ERASE.”
- The use of software RAID-based dynamic volumes is not recommended; use the array controller instead, it is more efficient.
- The largest disk that Windows Storage Server 2003 can accommodate from a storage system is 2 TB.
- It is not recommended to span array controllers with dynamic volumes.
- Use meaningful volume labels with the intended drive letter embedded in the volume label, if possible. For example, volume e: might be named “Disk E:.” Volume labels often serve as the only means of identification.
- Record all volume labels and drive letters in case of system Quick Restore. See “Managing Disks After Quick Restore” later in this chapter.

- When managing basic disks, only the last partition on the disk can be extended unless the disk is changed to dynamic.
- Basic disks can be converted to dynamic without bringing the system offline or loss of data, but the volume will be unavailable during the conversion.
- Basic disks can contain up to four primary partitions (or three primary partitions and one extended partition).
- Format drives with a 16 K allocation size for best support of snapshots, performance, and defragmentation.
- NTFS formatted drives are recommend since they provide the greatest level of support for snapshots, encryption, and compression.
- Only basic disks can be formatted as FAT or FAT32. Dynamic disks are not supported, nor can they be configured in a cluster.

Volumes Page

On the Volumes page, administrators can select to manage volumes, schedule defragmentation, and set or manage quotas. The Volumes page displays all volumes that are formatted NTFS on the system. It does not display the volume type (for example simple or spanned) nor volumes that are FAT32 or FAT. To display these types of volumes, click **Manage**. All volumes are displayed.

See the Managed Disks page to view a list of disks, and the volumes assigned to them. The drive letters for volumes that encompass multiple disks appear on multiple rows on the display.

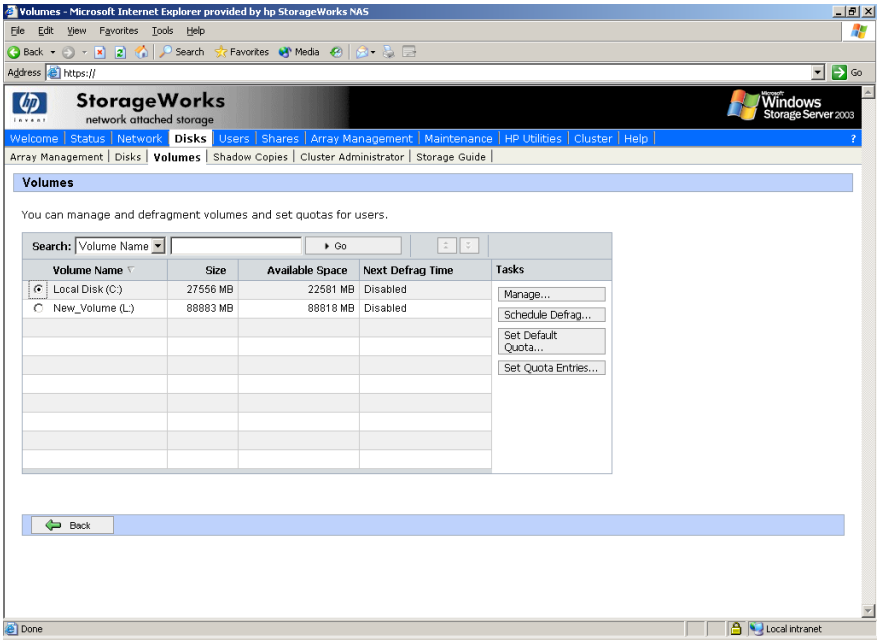

Figure 39: Volumes tab

Table 6: Volumes Page Object/Task Selector

Option	Task
Manage...	Select to display the Manage Volumes screen.
Schedule Defrag...	Select to schedule defragmentation for the selected volume.
Set Default Quota	Select to set quota limits to manage use of the volume. Settings on this page apply to new users and any users for whom user quota entries have not previously been set.
Set Quota Entries	Select to show a list of user quota entries. Then create a new quota entry, delete a quota entry, or view the properties of a quota entry.

Managing Volumes

To manage volumes on the server:

1. On the primary navigation bar, click **Disks**.
2. Click **Volumes**.
3. In the Tasks list, click **Manage**.

The Manage Volumes screen is displayed. The Manage Volumes page displays all volumes on the NAS device regardless of their format (NTFS, FAT, or FAT32). Do not tamper with the "Don't Erase" or the Local C: volume. These are reserved volumes and must be maintained as they exist.

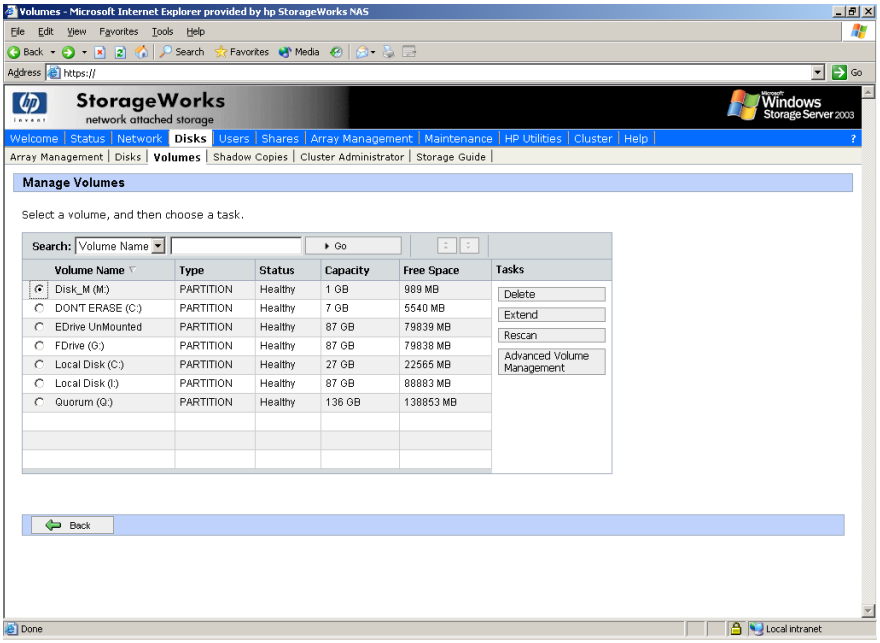

Figure 40: Manage Volumes screen

Table 7: Manage Volumes Options

Option	Task
Delete*	Select to delete the selected volume. This is data destructive and there is no recovery other than from tape.
Extend	Opens a page to extend a partition based on a basic disk or to extend dynamic based volumes.
Rescan	Select to detect a volume or partition added to the system or to update the size of a volume that has undergone expansion. The rescan is not synchronous and may require a browser refresh after the scan is initiated to display the new content.
Advanced Volume Management	Select to open the Windows Disk Manager and perform advanced volume management tasks.
* This task cannot be completed on a clustered resource.	

Dynamic Growth

Dynamic growth is a feature of the NAS server, which provides for growth of a volume or partition to meet expanding storage requirements without the need to take volumes offline or incur downtime. Growth may occur in three forms:

- Extend unallocated space from original LUNS.
- Alter LUNs to contain additional storage.
- Add new LUNS to the system. The additional space is then extended into through a variety of means depending on which type of disk structure is in use.

Expanding a LUN

Expanding an existing LUN is accomplished using the storage array configuration software applicable to the storage array in use. In the case of the MSA1000 controller, this is accomplished via the Array Configuration Utility presented on the Disk page. LUN expansion may occur in Disk Arrays where space is available. If insufficient space is available, additional physical disks may be added to the array dynamically.

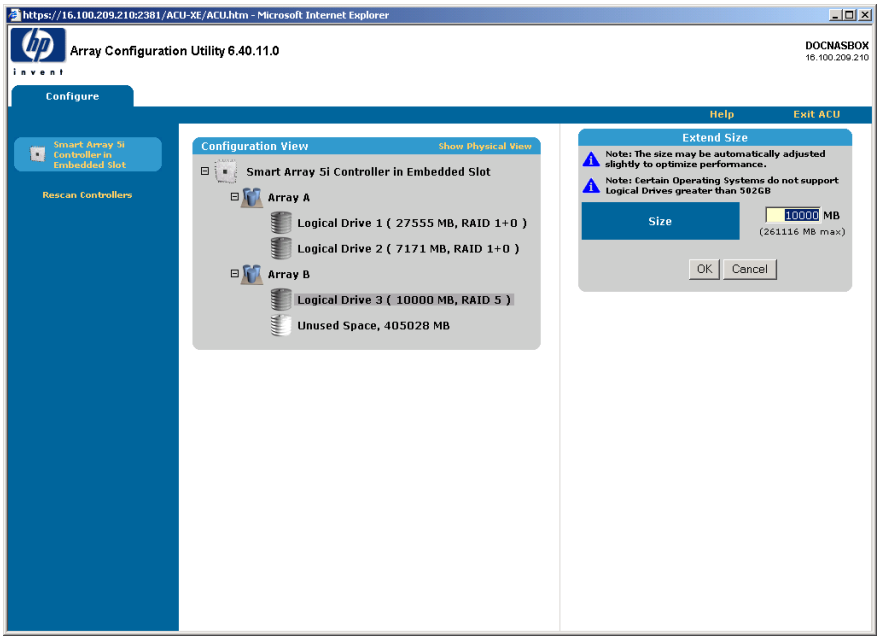

Figure 41: Expanding a LUN (MSA1000 only)

To extend a LUN where space is available in the array (MSA1000 only):

1. Click the **Disks** tab.
2. Click **Array Management**.
3. Click **Array Configuration Utility** and log in.
4. Select the appropriate array controller and the appropriate array that the logical drive is contained in.
5. Select the appropriate logical drive.
6. Select **Extend Size**.
7. Enter the total size of the logical drive in MB (not just the amount to be added) and click **OK**.

8. Click **Save** to update the configuration.
9. Close the ACU.

To extend a LUN where space is not available in the array (MSA1000 only):

1. Add an unassigned physical disk to the array using the ACU. If an unassigned physical disk is unavailable, add a new disk to the appropriate storage device and select **Refresh**.
2. To add an unassigned physical disk to the array use the following steps:
   a. Select the appropriate array controller and the appropriate array that the logical drive is contained in.
   b. Select **Expand Array**.
   c. Select the appropriate physical disk and click **OK**. The array is expanded.
3. Follow the instructions for extending a LUN.

Extending a partition on a basic disk

Partitions can be extended using either the WebUI extend function from the Managed Volumes page extend selection or by using the DiskPart command line utility. The Windows Disk Manager cannot extend basic disk partitions. To extend a partition using the WebUI follow the steps below:

1. Click the **Disks** tab.
2. Click the **Volumes** tab.
3. Click **Manage**.
4. Select the Volume to extend and click **Extend**.

Note: If you receive a message that there is not enough disk space to extend the volume, it is possible to convert the disk to dynamic, provided that there are other dynamic disks with space available and that the NAS device is not a node in a cluster. The volume can then be extended over a set of dynamic disks.

---

5. The page in Figure 42 is displayed. Enter in the amount to extend the partition.

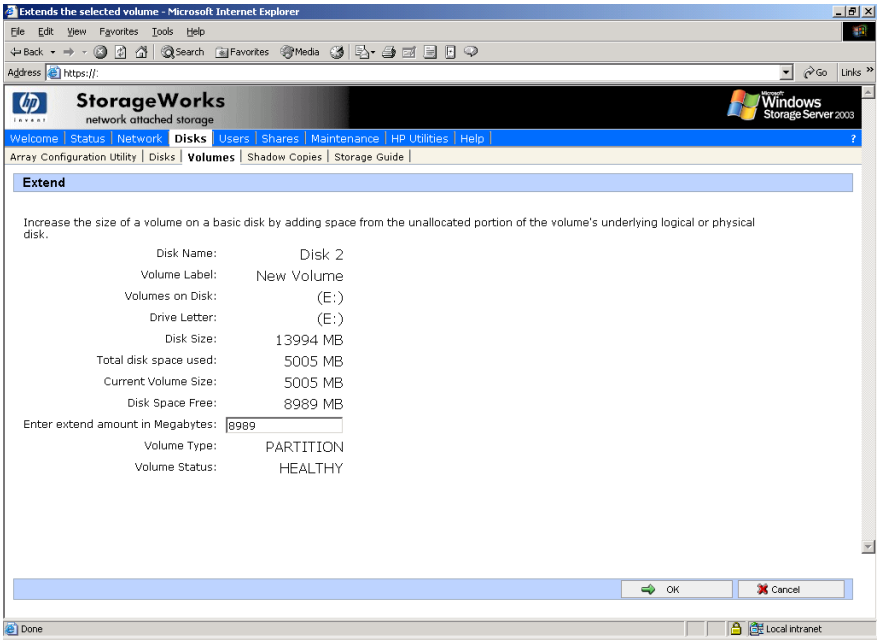

Figure 42: Extending a volume (basic disk)

6. Click **OK**.

Extending a Volume on Dynamic Disks (non-clustered systems only)

The WebUI allows for the extension of volumes based on a dynamic disk or a set of dynamic disks. To extend a volume perform the following steps:

1. Click the **Disks** tab.
2. Click the **Volumes** tab.
3. Click **Manage**.
4. Select the volume to extend and select **Extend**.

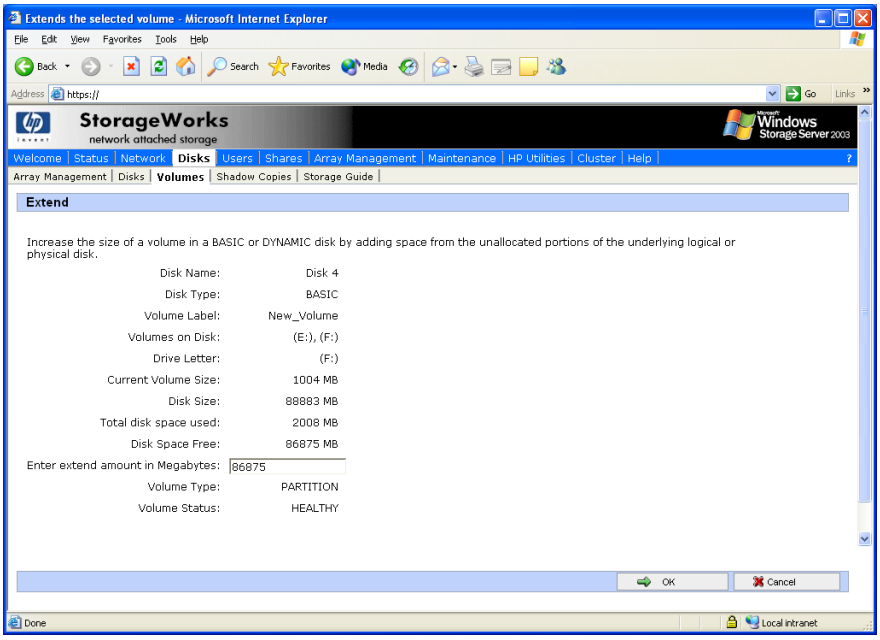

Figure 43: Extending a volume (dynamic disk)

5. Enter the amount to extend the volume. If no more space is available on the current dynamic disk, add an additional dynamic disk to the list of available disks and utilize space from it.
6. Click **OK**.

Extending using DiskPart

DiskPart may also be used to extend a partition or volume from the CMD prompt of the NAS operating system via Remote Desktop. Complete help is available from the Windows Storage Server 2003 desktop under **Start > Help and Support**. To use DiskPart follow the steps below:

Connect to the box through remote desktop, login, and select the command prompt icon.

1. Type `Diskpart`.
2. From the Diskpart prompt type the following commands:

- Type `list` to display all of the volumes
- Type `select [name of volume]` (for example `select Volume 4`) to work against a particular volume or partition.
- Type `Extend`. The volume is extended to the capacity of the underlying disk. To specify the amount to extend or to extend to another disk, type `extend [size=N] [disk=N]`

  Size is in MB.
- Type `exit` to exit the utility.

Scheduling Defragmentation

Defragmentation is the process of analyzing local volumes and consolidating fragmented files and folders so that each occupies a single, contiguous space on the volume. This allows the system to access files and folders and save new ones more efficiently. Because defragmentation consolidates files and folders, it also consolidates the free space on a volume. This reduces the likelihood that new files will be fragmented.

Defragmentation for a volume can be scheduled to occur automatically at convenient times. Defragmentation can also be done once, or on a recurring basis.

To schedule defragmentation for a volume:

1. On the primary navigation bar, choose **Disks**.
2. Click the **Volumes** tab.
3. Select the volume to schedule defragmentation.
4. In the Tasks list, choose **Schedule Defrag**.
5. On the **Manage the defragmentation schedule for [VolumeName]** page, select the **Schedule defragmentation for this volume** check box.
6. Select the frequency: Once, Weekly, or Monthly.
7. Use the remaining controls to specify when defragmentation will occur. The available controls change according to the frequency that is selected.
8. Click **OK**.

To disable defragmentation for a volume:

1. On the primary navigation bar, click **Disks**.
2. Click **Volumes**.
3. Select the volume to disable defragmentation.
4. In the Tasks list, choose **Schedule Defrag**.
5. On the **Manage the defragmentation schedule for [VolumeName]** page, clear the **Schedule defragmentation for this volume** check box.
6. Click **OK**.

Note: Scheduling defragmentation to run no later than a specific time prevents the defragmentation process from running later than that time. If the defragmentation process is running when the time is reached, the process is stopped. This setting is useful to ensure that the defragmentation process ends before the demand for server access is likely to increase.

If defragmenting volumes on which shadow copies are enabled, use a cluster (or allocation unit) size of 16 KB or larger during the format. Otherwise defragmentation registers as a change by the Shadow Copy process. This increase in the number of changes forces Shadow Copy to delete snapshots as the limit for the cache file is reached.

Caution: Allocation unit size cannot be altered without reformatting the drive. Data on a reformatted drive cannot be recovered.

Note: NTFS compression is supported only if the cluster size is 4 KB or smaller.

Managing Disks After Quick Restore

After a Quick Restore, drive letters may be assigned to the wrong volume. Windows Storage Server 2003 assigns drive letters after the restoration in the order of discovery. To help maintain drive letter information, placing the drive letter into the volume label is recommended. To change the drive letters to the appropriate ones, go into Disk Management and perform the following steps for each volume.

1. Right-click the on the volume that needs to be changed.
2. Select **Change Drive Letter and Paths**.
3. In the Change Drive Letter and Paths dialog box, select **Change**.

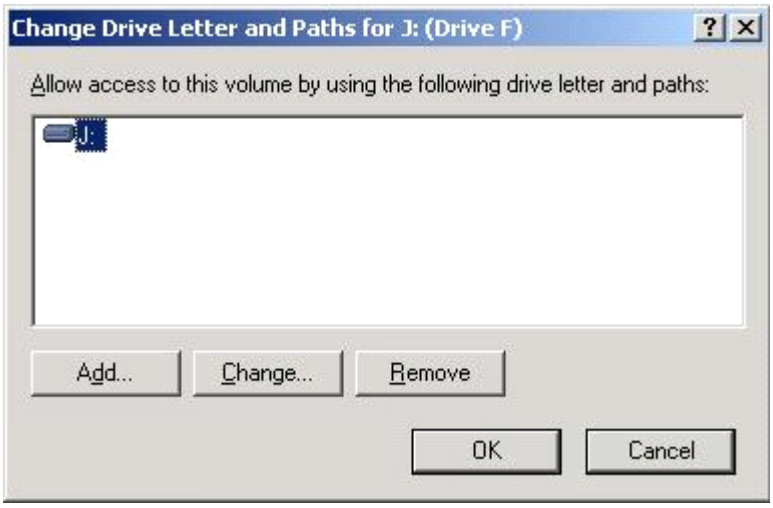

Figure 44: Change Drive Letter and Paths dialog box

4. In the Change Drive Letter or Path dialog box, select the appropriate drive letter and then click **OK**.

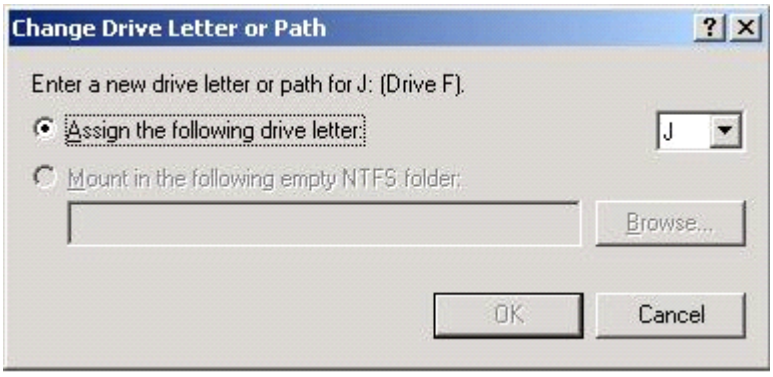

Figure 45: Enter new drive letter

5. Click **Yes** to confirm the drive letter change.

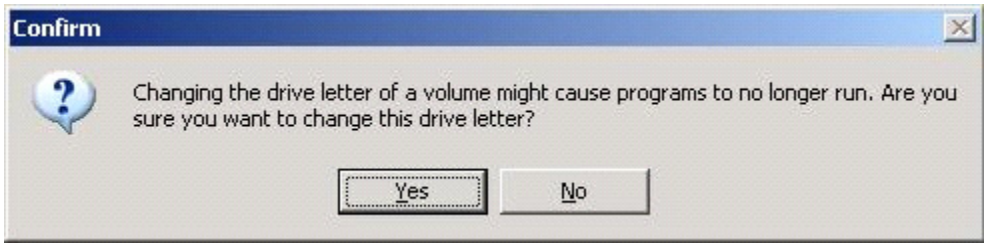

Figure 46: Confirm drive letter change

6. If the dialog box in Figure 47 is displayed, select **Yes** to continue. If the old drive letter needs to be reused, reboot the server after clicking **Yes**.

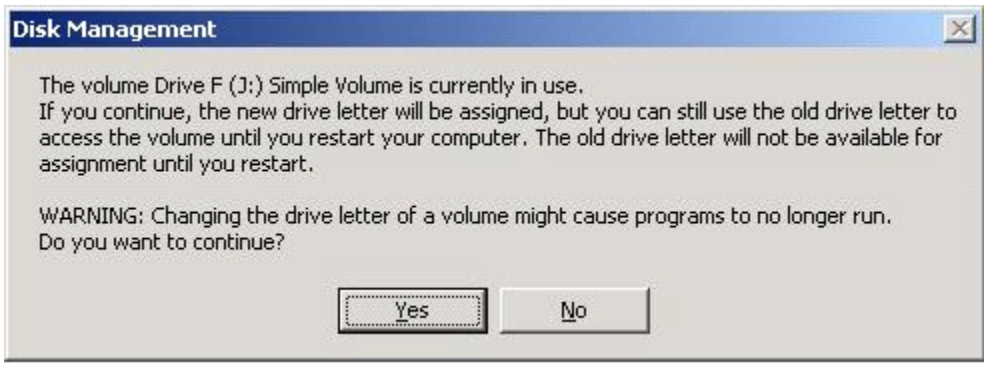

Figure 47: Disk Management warning

Disk Quotas

Disk quotas track and control disk space use in volumes.

Note: To limit the size of a folder or share, see "Directory Quotas" in Chapter 7.

Configure the volumes on the server to perform the following tasks:

- Prevent further disk space use and log an event when a user exceeds a specified disk space limit.
- Log an event when a user exceeds a specified disk space warning level.

When enabling disk quotas, it is possible to set both the disk quota limit and the disk quota warning level. The disk quota limit specifies the amount of disk space a user is allowed to use. The warning level specifies the point at which a user is nearing his or her quota limit. For example, a user's disk quota limit can be set to 50 megabytes (MB), and the disk quota warning level to 45 MB. In this case, the user can store no more than 50 MB on the volume. If the user stores more than 45 MB on the volume, the disk quota system logs a system event.

In addition, it is possible to specify that users can exceed their quota limit. Enabling quotas and not limiting disk space use is useful to still allow users access to a volume, but track disk space use on a per-user basis. It is also possible to specify whether or not to log an event when users exceed either their quota warning level or their quota limit.

When enabling disk quotas for a volume, volume usage is automatically tracked from that point forward, but existing volume users have no disk quotas applied to them. Apply disk quotas to existing volume users by adding new quota entries on the Quota Entries page.

Enabling Quota Management

When enabling disk quotas on a volume, every user's disk volume usage is monitored and treated differently, depending on the quota management settings for the specific user.

To enable quota management on a volume:

1. On the primary navigation bar, click **Disks**.
2. Click **Volumes**.
3. Select the volume to manage.
4. In the Tasks list, click **Set Default Quota**.
5. On the Default Quota for volume page, select **Use quota limits to manage use of the volume**.
6. If desired, select **Deny disk space to users exceeding quota limit** to enable that restriction.
7. Specify the default quota limit and warning level for new users on this volume.
8. Specify which quota events should be logged.
9. Click **OK**.

Note: When enabling disk quotas on a volume, any users with write access to the volume who have not exceeded their quota limit can store data on the volume. The first time a user writes data to a quota-enabled volume, default values for disk space limit and warning level are automatically assigned by the quota system.

To disable quota management on a volume:

1. On the primary navigation bar, click **Disks**.
2. Click **Volumes**.
3. Select the volume to manage.
4. In the Tasks list, click **Set Default Quota**.
5. On the Default Quota for (volume) page, clear the check box to **Use quota limits to manage use of the volume**.
6. Click **OK**.

Setting User Quota Entries

The Set User Quotas page allows the administrator to set, delete, or change disk quotas for any user on the server. To set or change quota entries on the server:

1. On the primary navigation bar, click **Disks**.
2. Click **Volumes**.
3. Select the volume to manage.
4. From the Tasks list, click **Set Quota Entries**.

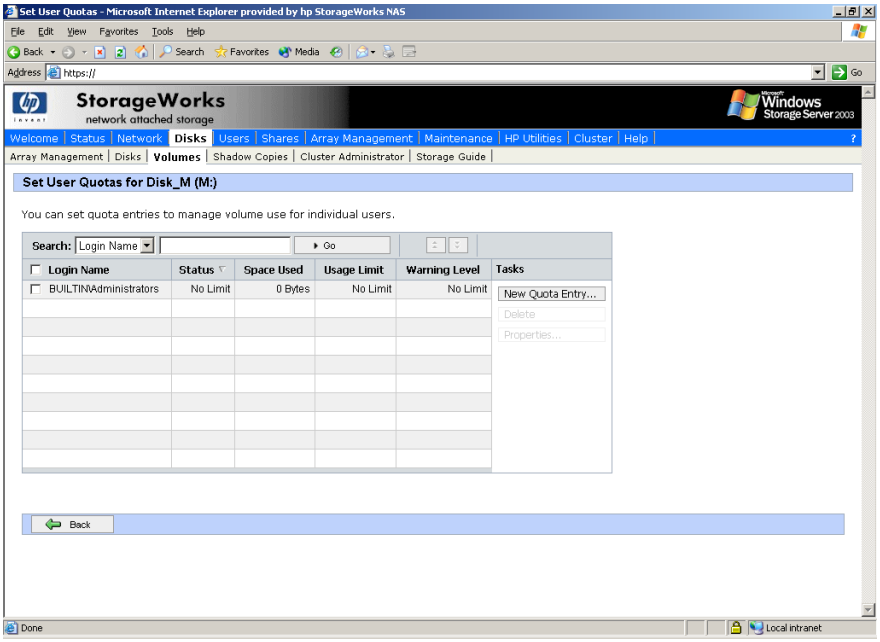

Figure 48: Setting user quotas

To create a new user quota entry:

1. Click **New Quota Entry**.
2. Select a user.
3. Set the limit.
4. Set the warning level.
5. Click **OK**.

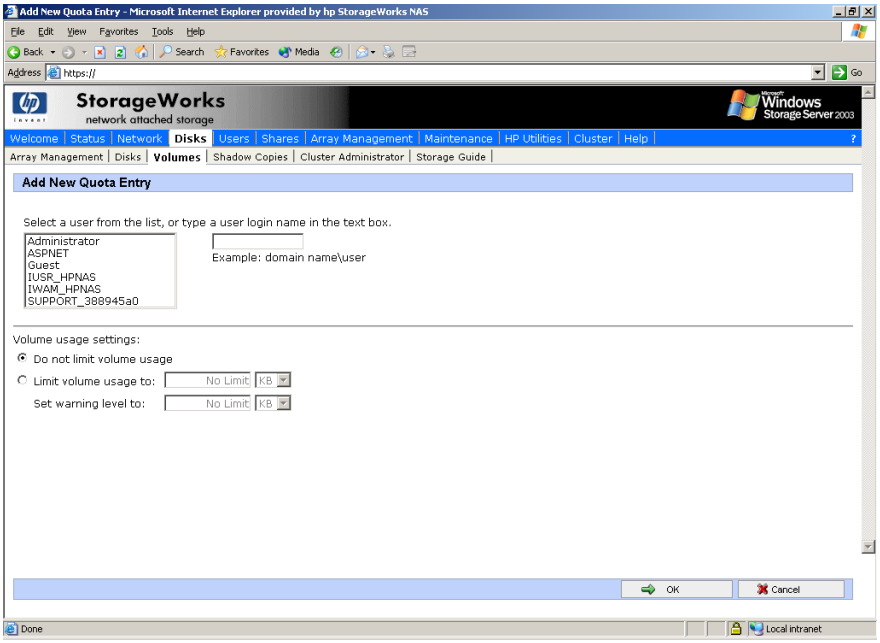

Figure 49: Add new quota entry

To change a quota entry:

1. Select the quota to change.
2. Click **Properties**.
3. Change the limit.
4. Change the warning level.
5. Click **OK**.

To delete a quota entry:

1. Select the quota to change.
2. Click **Delete**.

DiskPart

DiskPart.exe is a text-mode command interpreter that enables the administrator to manage disks, partitions, or volumes.

When using the list commands, an asterisk (*) appears next to the object with focus. Select an object by its number or drive letter, such as disk 0, partition 1, volume 3, or volume C.

When selecting an object, the focus remains on that object until a different object is selected. For example, if the focus is set on disk 0 and volume 8 on disk 2 is selected, the focus shifts from disk 0 to disk 2, volume 8. Some commands automatically change the focus. For example, when creating a new partition, the focus automatically switches to the new partition.

Focus can only be given to a partition on the selected disk. When a partition has focus, the related volume (if any) also has focus. When a volume has focus, the related disk and partition also have focus if the volume maps to a single specific partition. If this is not the case, focus on the disk and partition is lost.

Table 8: Common DiskPart Commands

Command	Description
`add disk`	Mirrors the simple volume with focus to the specified disk.
`assign`	Assigns a drive letter or mount point to the volume with focus.
`convert basic`	Converts an empty dynamic disk to a basic disk.
`convert dynamic`	Converts a basic disk into a dynamic disk. Any existing partitions on the disk become simple volumes.
`create volume simple`	Creates a simple volume. After creating the volume, the focus automatically shifts to the new volume.
`exit`	Exits the DiskPart command interpreter.
`help`	Displays a list of the available commands.
`list disk`	Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table. The disk marked with an asterisk (*) has focus.
`list partition`	Displays the partitions listed in the partition table of the current disk. On dynamic disks these partitions may not correspond to the dynamic volumes on the disk. This discrepancy occurs because dynamic disks contain entries in the partition table for the system volume or boot volume (if present on the disk). They also contain a partition that occupies the remainder of the disk in order to reserve the space for use by dynamic volumes.
`list volume`	Displays a list of basic and dynamic volumes on all disks.
`rem`	Provides a way to add comments to a script.
`retain`	Prepares an existing dynamic simple volume to be used as a boot or system volume.
`select disk`	Selects the specified disk and shifts the focus to it.

For a complete list of DiskPart commands, go to the Windows Storage Server 2003 Desktop on the NAS device via Remote Desktop and select **Start >Help and Support**, search on DiskPart.

Example of using DiskPart

The following example shows how to configure a volume on the NAS server.

In the cmd window, type:

```
c:\>diskpart
DISKPART>Rescan
DISKPART>select disk 2
DISKPART>convert dynamic
DISKPART>REM Create a simple volume
DISKPART>create volume simple size=4000
DISKPART> REM Assign drive letter E: to the volume
DISKPART>assign letter=E
DISKPART>list vol
DISKPART>Exit
```

Shadow Copies

5

Overview

Note: The NAS 4000s and 9000s servers can be deployed in a clustered as well as a non-clustered configuration. This chapters discusses using Shadow Copies in a non-clustered environment. Please review the Cluster Administration chapter of this guide for additional information regarding Shadow Copies in a cluster.

The Volume Shadow Copy Service provides an infrastructure for creating point-in-time snapshots (shadow copies) of volumes. Shadow Copy supports 64 shadow copies per volume.

A shadow copy contains previous versions of the files or folders contained on a volume at a specific point in time. While the Shadow Copy mechanism is managed at the server (see the "Managing Shadow Copy" section in this chapter), previous versions of files and folders are only available over the network from clients and are seen on a per folder or file level and not as an entire volume.

The Shadow Copy feature works at the block level. As changes are made to the file system, the Shadow Copy Service copies out the original blocks to a special cache file, to maintain a consistent view of the file at a particular point in time. Since the snapshot only contains a subset of the original blocks, the cache file is typically smaller than the original volume. In the snapshot's original form, it takes up no space since blocks are not moved until an update to the disk occurs.

By using shadow copies, a NAS server can maintain a set of previous versions of all files on the selected volumes. End users access the file or folder by using a separate client add-on program, which enables them to view the file in Windows Explorer. Accessing previous versions of files, or shadow copies, enables users to:

- Recover files that were accidentally deleted. Previous versions can be opened and copied to a safe location.
- Recover from accidentally overwriting a file. A previous version of that file can be accessed.
- Compare several versions of a file while working. Use previous versions to compare changes between two versions of a file.

Shadow copies cannot replace the current backup, archive, or business recovery system, but they can help to simplify restore procedures. Since a snapshot only contains a portion of the original data blocks, shadow copies can not protect against data loss due to media failures. However the strength of snapshots is the ability to instantly recover data from shadow copies, reducing the number of times needed to restore data from tape.

Shadow Copy Planning

Before setup is initiated on the server and the client interface is made available to end users, consider the following:

- From what volume will shadow copies be taken?
- How much disk space should be allocated for shadow copies?
- Will separate disks be used to store shadow copies?
- How frequently will shadow copies be made?

Identifying the Volume

Shadow copies are taken for a complete volume, but not for a specific directory. Shadow copies work best when the server stores user files, such as documents, spreadsheets, presentations, graphics, or database files.

Note: Shadow copies should not be used to provide access to previous versions of application or e-mail databases.

Shadow copies are designed for volumes that store user data such as home directories and My Documents folders that are redirected by using Group Policy or other shared folders in which users store data.

Shadow copies work with compressed or encrypted files and retain whatever permissions were set on the files when the shadow copies were taken. For example, if a user is denied permission to read a file, that user would not be able to restore a previous version of the file, or be able to read the file after it has been restored.

Although shadow copies are taken for an entire volume, users must use shared folders to access shadow copies. Administrators on the local server must also specify the *servername**sharename* path to access shadow copies. If administrators or end users want to access a previous version of a file that does not reside in a shared folder, the administrator must first share the folder.

Note: Shadow copies are available only on NTFS, not FAT or FAT32 volumes.

Files or folders that are recorded by using Shadow Copy appear static, even though the original data is changing.

Allocating Disk Space

When shadow copies are enabled on a volume, the maximum amount of volume space to be used for the shadow copies can be specified. The default limit is 10 percent of the source volume (the volume being copied). The limit for volumes in which users frequently change files should be increased. Also, note that setting the limit too low causes the oldest shadow copies to be deleted frequently, which defeats the purpose of shadow copies and frustrates users.

If the frequency of changes to each file is greater than the amount of space allocated to storing shadow copies, then no shadow copy is created. Therefore, administrators should carefully consider the amount of disk space they want to set aside for shadow copies, and keep in mind

user expectations of how many versions they will want to have available. End users might expect only a single shadow copy to be available, or they might expect three days or three weeks worth of shadow copies. The more shadow copies users expect, the more storage space administrators must allocate for storing them.

Setting the limit too low also affects Backup and other backup programs that use shadow copy technology because these programs are also limited to using the amount of disk space specified by administrators.

Note: Regardless of the volume space that is allocated for shadow copies, there is a maximum of 64 shadow copies for any volume. When the 65th shadow copy is taken, the oldest shadow copy is purged.

The minimum amount of storage space that can be specified is 100 megabytes (MB). The default storage size is 10% of the source volume (the volume being copied). If the shadow copies are stored on a separate volume, change the default to reflect the space available on the storage volume instead of the source volume. Remember that when the storage limit is reached, older versions of the shadow copies are deleted and cannot be restored.

When determining the amount of space to allocate for storing shadow copies, consider both the number and size of files that are being copied, as well as the frequency of changes between copies. For example, 100 files that only change monthly require less storage space than 10 files that change daily.

To change the storage volume, shadow copies must be deleted. The existing file change history that is kept on the original storage volume is lost. To avoid this problem, verify that the storage volume that is initially selected is large enough.

When using a basic disk as a storage area for shadow copies and converting the disk into a dynamic disk, it is important to take the following precaution to avoid data loss:

- If the disk is a non-boot volume and is a different volume from where the original files reside, first dismount and take offline the volume containing the original files before converting the disk containing shadow copies to a dynamic disk.
- The volume containing the original files must be brought back online within 20 minutes, otherwise, the data stored in the existing shadow copies is lost.
- If the shadow copies are located on a boot volume, the disk to can be converted to dynamic without losing shadow copies.

Note: Use the mountvol command with the /p option to dismount the volume and take it offline. Mount the volume and bring it online using the mountvol command or the Disk Management snap-in.

Identifying the Storage Area

To store the shadow copies of another volume on the same file server, a volume can be dedicated on separate disks. For example, if user files are stored on *H:*\, another volume such as *S:*\ can be used to store the shadow copies. Using a separate volume on separate disks provides better performance and is recommended for heavily used NAS devices.

If a separate volume will be used for the storage area (where shadow copies are stored), the maximum size should be changed to **No Limit** to reflect the space available on the storage area volume instead of the source volume (where the user files are stored).

Disk space for shadow copies can be allocated on either the same volume as the source files or a different volume. There is, however, a trade-off between ease of use and maintenance versus performance and reliability that the system administrator must consider.

For example, by keeping the shadow copy on the same volume, although there is a potential gain in ease of setup and maintenance, there may be a reduction in performance and reliability.

Caution: If shadow copies are stored on the same volume as the user files, note that a burst of disk input/output (I/O) can cause all shadow copies to be deleted. If the sudden deletion of shadow copies is unacceptable to administrators or end users, it is best to use a separate volume on separate disks to store shadow copies.

Determining Creation Frequency

The more frequently shadow copies are created, the more likely that end users will get the version that they want. However, with a maximum of 64 shadow copies per volume, there is a trade-off between the frequency of making shadow copies and the amount of time that the earlier files will be available.

By default, the NAS server will create shadow copies at 0700 and 1200, Monday through Friday when the feature is enabled for a volume. However, these settings are easily modified by the administrator so that the shadow copy schedule can better accommodate end user needs. To modify these schedules see the section on “Shadow Copy Schedules” documented later in this chapter.

Note: The more shadow copies are created, the more disk space the shadow copies can consume, especially if files change frequently.

Shadow Copies and Drive Defragmentation

When running Disk Defragmenter on a volume with shadow copies activated, all or some of the shadow copies may be lost, starting with the oldest shadow copies.

If defragmenting volumes on which shadow copies are enabled, use a cluster (or allocation unit) size of 16 KB or larger. Utilizing this allocation unit size reduces the number of copy outs occurring on the snapshot. Otherwise the number of changes caused by the defragmentation process can cause shadow copies to be deleted faster than expected. Note, however, that NTFS compression is supported only if the cluster size is 4 KB or smaller.

Note: To check the cluster size of a volume, use the `fsutil fsinfo ntfsinfo` command. To change the cluster size on a volume that contains data, backup the data on the volume, reformat it using the new cluster size, and then restore the data.

Mounted Drives

A mounted drive is a local volume attached to an empty folder (called a mount point) on an NTFS volume. When enabling shadow copies on a volume that contains mounted drives, the mounted drives are not included when shadow copies are taken. In addition, if a mounted drive is shared and shadow copies are enabled on it, users cannot access the shadow copies if they traverse from the host volume (where the mount point is stored) to the mounted drive.

For example, assume there is a folder *E:\data\users*, and the *Users* folder is a mount point for *F:*. If shadow copies are enabled on both *E:* and *F:*, *E:\data* is shared as *server1**data*, and *E:\data\users* is shared as *server1**users*. In this example, users can access previous versions of *server1**data* and *server1**users* but not *server1**data**users*.

Managing Shadow Copies

From the **WebUI Welcome** screen, click **Disks**, then **Shadow Copies** to display the Shadow Copies screen.

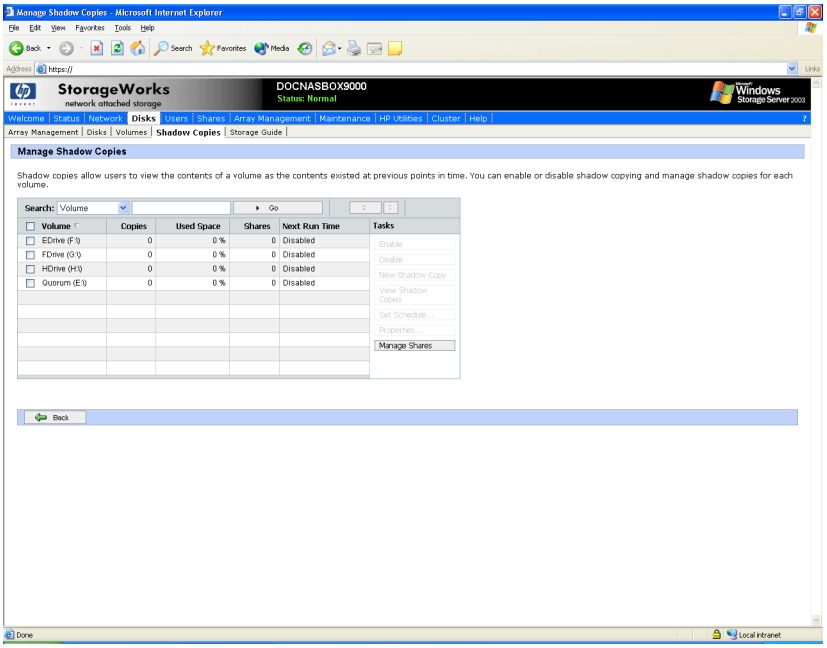

Figure 50: Shadow Copies screen

Table 9: Shadow Copies Fields

Field	Description
Volume	Lists all volumes of the server on which the Shadow Copies service can be used. Only NTFS file system data volumes that are physically located on the server can support shadow copies. To manage shadow copies on a volume, select the check box next to the volume name, and then choose a task from the Tasks list.
Copies	Lists the number of shadow copies on the volume.
Used Space	Lists the total disk space that is used by the shadow copies on the volume.
Shares	Lists the number of shared folders that reside on the volume. This information can help determine whether to enable shadow copies on a volume. A greater number of shared folders on a volume increases the likelihood that users might need access to previous versions of their data.
Next Run Time	If the Shadow Copies service is enabled on the volume, this column lists the time and date the next shadow copy will be created. Otherwise, it displays Disabled.

Table 10: Shadow Copies Tasks

Task	Description
Enable	Click to enable Shadow Copies on the selected volume.
Disable	Click to enable Shadow Copies on the selected volume.
New Shadow Copy	Click to immediately create a new shadow copy on the selected volume.
View Shadow Copies	Click to view a list of shadow copies on the selected volume.
Set Schedule	Click to set the time and frequency of shadow copies.
Properties...	Click to view the shadow copy properties of the selected volume, including location and size of the cache file.
Manage Shares	Click to go to the Shared Folders screen.

The Shadow Copy Cache File

The default shadow copy settings allocate 10% of the source volume being copied (with a minimum of 100 MB), and store the shadow copies on the same volume as the original volume. See Figure 51. The cache file is located in a hidden protected directory entitled "System Volume Information" off of the root of each volume for which Shadow Copy is enabled.

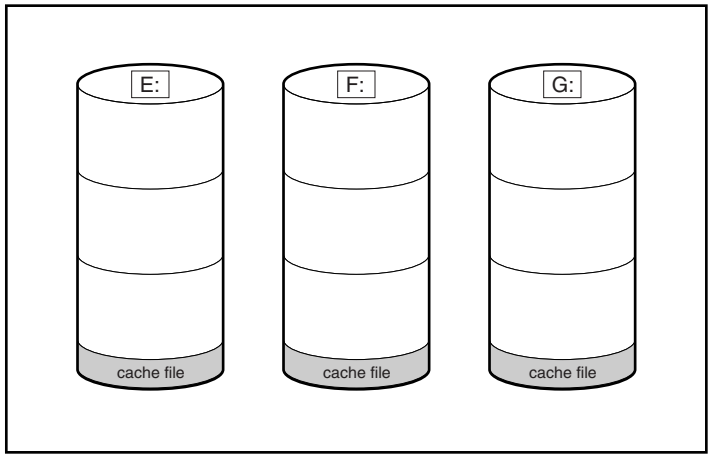

Figure 51: Shadow copies stored on source volume

As mentioned previously, the cache file location can be altered to reside on a dedicated volume separate from the volumes containing files shares. See Figure 52.

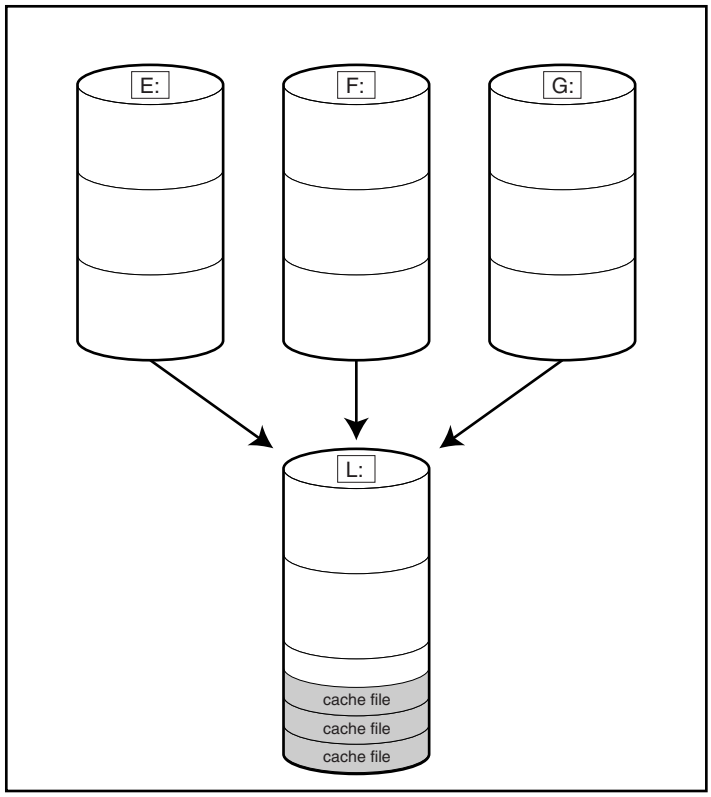

Figure 52: Shadow copies stored on separate volume

The main advantage to storing shadow copies on a separate volume is ease of management and performance. Shadow copies on a source volume must be continually monitored and can consume space designated for file sharing. Setting the limit too high takes up valuable storage space. Setting the limit too low can cause shadow copies to be purged too soon, or not created at all. By storing shadow copies on a separate volume space for Shadow Copies may be manage separately, limits can generally be set higher, or set to No Limit. See the properties tab of the shadow copy page for a volume to alter the cache file location, covered later in this chapter.

Caution: If the data on the separate volume L: is lost, the shadow copies cannot be recovered.

Enabling and Creating Shadow Copies

Enabling the Shadow Copies service for a volume or creating a shadow copy can be done directly from the Manage Shadow Copies page.

Enabling shadow copies on a volume automatically results in several actions:

- Creates a shadow copy of the selected volume
- Sets the maximum storage space for the shadow copies
- Schedules shadow copies to be made at 7 A.M. and 12 noon on weekdays.

Note: Creating a shadow copy only makes one copy of the volume; it does not create a schedule.

To enable shadow copies on a volume:

1. On the primary navigation bar, click **Disks**.
2. Click the **Shadow Copies** tab.
3. On the Manage Shadow Copies page, select one or more volumes to enable the Shadow Copies service on.

Note: After the first shadow copy is created, it cannot be relocated. Relocate the cache file by altering the cache file location under Properties prior to enabling shadow copy. See "Viewing Shadow Copy Properties" in this chapter.

4. Click **Enable**.

To create a shadow copy on a volume:

1. On the primary navigation bar, click **Disks**.
2. Click the **Shadow Copies** tab.
3. On the Manage Shadow Copies page, select one or more volumes on which to create the shadow copies.
4. Click **New Shadow Copy**.

Viewing a List of Shadow Copies

To view a list of shadow copies on a volume:

1. On the primary navigation bar, click **Disks**.
2. Click the **Shadow Copies** tab.
3. On the Manage Shadow Copies page, select the volume to view.
4. On the Tasks list, click **View Shadow Copies**.

   All shadow copies are listed, sorted by the date and time they were created.

Note: It is also possible to create new shadow copies or delete shadow copies from this page.

Set Schedules

Shadow Copy schedules control how frequently shadow copies of a volume are made. There are a number of factors that can help determine the most effective shadow copy schedule for an organization. These include the work habits and locations of the users. For example, if users do not all live in the same time zone, or they work on different schedules, it is possible to adjust the daily shadow-copy schedule to allow for these differences.

It is recommended that shadow copies be scheduled not more frequently than once per hour.

Scheduling Shadow Copies

When the Shadow Copies service is enabled on a volume, it automatically schedules shadow copies to be made each weekday at 7 A.M. and 12 noon.

To add or change a shadow copy schedule for a volume:

1. On the primary navigation bar, click **Disks**.
2. Click **Shadow Copies**.
3. Select the volume.
4. In the Tasks list, click **Set Schedule**.
5. On the Shadow Copy Schedules page, click **New**.
6. Select a frequency: Once, Daily, Weekly, or Monthly.
7. Use the remaining controls to specify the recurrence pattern and the starting date and time. The available controls change according to the frequency selected.
8. Click **OK**.

Deleting a Shadow Copy Schedule

To delete a shadow copy schedule on a volume:

1. On the primary navigation bar, click **Disks**.
2. Click the **Shadow Copies** tab.
3. Select the volume on which to delete a shadow copy schedule.
4. In the Tasks list, click **Set Schedule**.
5. On the Manage Shadow Copy Schedules screen, select the schedule to be deleted, and click **Delete**.
6. Click **OK** to confirm the deletion or **Cancel** to retain the copy.

Note: When deleting a shadow copy schedule, that action has no effect on existing shadow copies. To remove schedules and all shadow copies in one action, from the Manage Shadow Copies page, choose Disable from the Tasks list.

Viewing Shadow Copy Properties

To view shadow copy properties on a volume:

1. On the primary navigation bar, click **Disks**.
2. Click the **Shadow Copies** tab.

3. On the Manage Shadow Copies page, select the volume on which to view shadow copy properties.
4. On the Tasks list, click **Properties**.

   The Shadow Copy Properties screen, as shown in Figure 53, lists the number of copies, the date and time the most recent shadow copy was made, and the maximum size setting.

   Change the maximum size limit for all shadow copies, or choose **No limit**.

   For volumes where shadow copies do not exist currently, it is possible to change the location of the cache file. See "The Shadow Copy Cache File" earlier in this chapter. The list of available disks and the space available on each is presented at the bottom of the page. Managing the cache files on a separate disk is recommended.

Note: If shadow copies have already been enabled, the cache file location is grayed out. To change this location after shadow copies have been enabled, all shadow copies must be deleted and cannot be recovered. Remember enabling Shadow Copies creates a Shadow Copy by default.

5. Click **OK** to save changes, or click **Cancel** to discard changes.

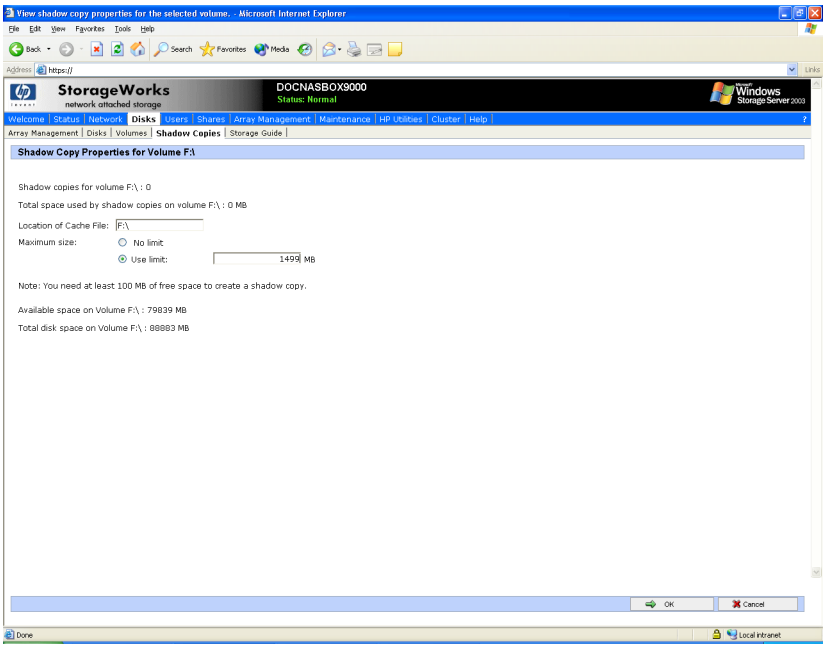

Figure 53: Shadow Copies properties screen

Caution: Use caution when reducing the size limit for all shadow copies. When the size is set to less than the total size currently used for all shadow copies, enough shadow copies are deleted to reduce the total size to the new limit. A shadow copy cannot be recovered after it has been deleted.

Disabling Shadow Copies

When shadow copies are disabled on a volume, all existing shadow copies on the volume are deleted as well as the schedule for making new shadow copies.

To disable shadow copies on a volume:

1. On the primary navigation bar, click **Disks**.
2. Click the **Shadow Copies** tab.
3. On the Manage Shadow Copies page, select one or more volumes on which to disable shadow copies.
4. In the Tasks list, click **Disable**.

   The Disable Shadow Copies page identifies the volume for which shadow copies will be disabled.
5. Click **OK** to delete all existing shadow copies and settings for the volume.

Caution: When the Shadow Copies service is disabled, all shadow copies on the selected volumes are deleted. Once deleted, shadow copies cannot be restored.

Managing Shadow Copies from the NAS Desktop

As an alternative to managing Shadow Copies via the WebUI, the NAS Desktop may be accessed via Remote Desktop.

To access Shadow Copies from the NAS Desktop:

1. From the WebUI select **Remote Desktop** from the Maintenance tab.
2. Click **My Computer**.
3. Select the volume.
4. Right-click the volume name and select **Properties**.
5. Click the **Shadow Copies** tab.

The user interface provides the same functionality found in the WebUI but in Win32 form. See Figure 54.

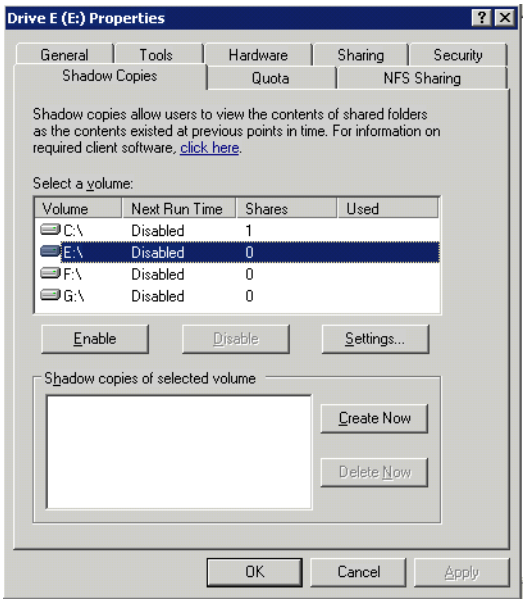

Figure 54: Accessing Shadow Copies from My Computer

Shadow Copies for Shared Folders

Shadow Copies are accessed over the network by supported clients and protocols. There are two sets of supported protocols, SMB and NFS. All other protocols are not supported; this would include HTTP, FTP, AppleTalk, and NetWare Shares. For SMB support a client side application denoted as Shadow Copies for Shared Folders is required. The client side application is currently only available for Windows XP and Windows 2000 SP3+. The application can be downloaded by clicking on **Shadow Copy Client** from the Microsoft website:

http://www.microsoft.com/windowsserver2003/downloads/featurepacks/default.mspx

No additional software is required to enable UNIX users to independently retrieve previous versions of files stored on NFS shares.

Note: Shadow Copies for Shared Folders supports retrieval only of shadow copies of network shares. It does not support retrieval of shadow copies of local folders.

Note: Shadow Copies for Shared Folders clients are not available for HTTP, FTP, AppleTalk, or NetWare shares. Consequently, users of these protocols cannot use Shadow Copies for Shared Folders to independently retrieve previous versions of their files. However, administrators can take advantage of Shadow Copies for Shared Folders to restore files on behalf of these users.

SMB Shadow Copies

Windows users can independently access previous versions of files stored on SMB shares via the Shadow Copies for Shared Folders client. After the Shadow Copies for Shared Folders client is installed on the user's computer, the user can access shadow copies for a share by right-clicking on the share to open its Properties dialog, selecting the Previous Versions tab, and then selecting the desired shadow copy. Users can view, copy, and restore all available shadow copies.

Shadow Copies for Shared Folders preserves the permissions set in the access control list (ACL) of the original folders and files. Consequently, users can only access shadow copies for shares to which they have access. In other words, if a user does not have access to a share, he also does not have access to the share's shadow copies.

The Shadow Copies of Shared Folders client pack installs a **Previous Versions** tab in the **Properties** dialog box of files and folders on network shares.

Users access shadow copies with Windows Explorer by selecting **View**, **Copy**, or **Restore**, from the **Previous Versions** tab. See Figure 55. Both individual files and folders may be restored.

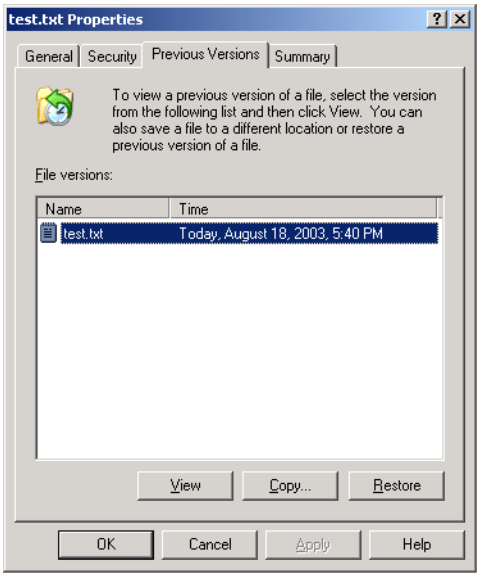

Figure 55: Client GUI

When users view a network folder hosted on the NAS device for which shadow copies are enabled, old versions (prior to the snapshot) of a file or directory are available. Viewing the properties of the file or folder presents users with the folder or file history-a list of read-only, point-in-time copies of the file or folder contents that users can then open and explore like any other file or folder. Users can view files in the folder history, copy files from the folder history, and so on.

NFS Shadow Copies

UNIX users can independently access previous versions of files stored on NFS shares via the NFS client; no additional software is required. Server for NFS exposes each of a share's available shadow copies as a pseudo-subdirectory of the share. Each of these pseudo-subdirectories is displayed in exactly the same way as a regular subdirectory is displayed.

The name of each pseudo-subdirectory reflects the creation time of the shadow copy, using the format.@GMT-YYYY.MM.DD-HH:MM:SS. Note that, to prevent common tools from needlessly enumerating the pseudo-subdirectories, the name of each pseudo-subdirectory begins with the dot character, thus rendering it hidden.

The following example shows an NFS share named “NFSShare” with three shadow copies, taken on April 27, 28, and 29 of 2003 at 4 a.m.

NFSShare

.@GMT-2003.04.27-04:00:00

.@GMT-2003.04.28-04:00:00

.@GMT-2003.04.29-04:00:00

Access to NFS shadow copy pseudo-subdirectories is governed by normal access-control mechanisms using the permissions stored in the file system. Users can access only those shadow copies to which they have read access at the time the shadow copy is taken. To prevent users from modifying shadow copies, all pseudo-subdirectories are marked read-only, regardless of the user's ownership or access rights, or the permissions set on the original files.

Server for NFS periodically polls the system for the arrival or removal of shadow copies and updates the root directory view accordingly. Clients then capture the updated view the next time they issue a directory read on the root of the share.

Recovery of Files or Folders

There are three common situations that may require recovery of files or folders:

- Accidental file deletion, the most common situation.
- Accidental file replacement, which may occur if a user selects Save instead of Save As.
- File corruption.

It is possible to recover from all of these scenarios by accessing shadow copies. There are separate steps for accessing a file compared to accessing a folder.

As documented previously, the use of the snapshots are from the network and are based on shares created on the NAS server.

Recovering a Deleted File or Folder

To recover a deleted file or folder within a folder:

1. Navigate to the folder where the deleted file was stored.
2. Position the cursor over a blank space in the folder. If the cursor hovers over a file, that file will be selected.
3. Right-click the mouse and select **Properties** from the bottom of the menu. Select the **Previous Versions** tab.
4. Select the version of the folder that contains the file before it was deleted, and then click **View**.
5. View the folder and select the file or folder to recover. The view may be navigated multiple folders deep.
6. Select restore to restore the file or folder to its original location. Selecting copy will allow the placement of the file or folder to a new location.

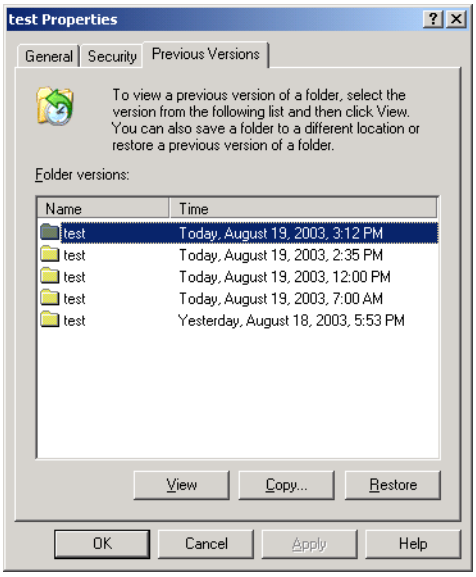

Figure 56: Recovering a deleted file or folder

Recovering an Overwritten or Corrupted File

Recovering an overwritten or corrupted file is easier than recovering a deleted file because the file itself can be right-clicked instead of the folder. To recover an overwritten or corrupted file use the following procedure:

1. Right-click the overwritten or corrupted file and click **Properties**.
2. Select **Previous Versions**.
3. To view the old version, click **View**. To copy the old version to another location, click **Copy...** to replace the current version with the older version, click **Restore**.

Recovering a Folder

To recover a folder use the following procedure:

1. Position the cursor so that it is over a blank space in the folder that will be recovered. If the cursor hovers over a file, that file will be selected.
2. Right-click the mouse, select **Properties** from the bottom of the menu, then click the **Previous Versions** tab.
3. Choose either **Copy** or **Restore**.
4. Choosing **Restore** enables the user to recover everything in that folder as well as all subfolders. Selecting **Restore** will not delete any files.

Backup and Shadow Copies

As mentioned previously, Shadow Copies are only available on the network via the client application and only at a file or folder level as opposed to the entire volume. Hence the standard backup associated with a volume backup will not work to back up the previous versions of the file system. To answer this particular issue, Shadow Copies are available for back up in two situations. If the backup software in question supports the use of Shadow Copies and can communicate with underlying block device, it is supported and the previous version of the file system will be listed in the backup application as a complete file system snapshot. Lastly, if the built in backup application NTbackup is utilized, the backup software forces a snapshot and then uses the snapshot as the means for backup. The user is unaware of this activity and it is not self evident although it does address the issue of open files.

User and Group Management

6

There are two system environments for users and groups: workgroup and domain. Because users and groups in a domain environment are managed through standard Windows NT or Windows Storage Server 2003 domain administration methods, this document discusses only local users and groups, which are stored and managed on the NAS device. For information on managing users and groups on a domain, refer to the domain documentation available on the Microsoft website.

Domain Compared to Workgroup Environments

NAS server devices can be deployed in workgroup or domain environments. When in a domain environment, the server is a member of the domain. The domain controller is a repository of accounts and account access for the NAS server. Client machines are also members of the domain, and users log on to the domain through their Windows clients. The domain controller also administers user accounts and appropriate access levels to resources that are a part of the domain.

In a CIFS/SMB environment, when mapping a network drive or a client machine, a user sends a logon credential to the server. This credential includes the username, password, and if appropriate, domain information. Using the credential, the server authenticates and provides the corresponding access to the user.

When a NAS server is deployed into a workgroup environment, all user and group account access permissions to file resources are stored locally on the server.

By contrast, when a NAS server is deployed into a domain environment it uses the account database from the domain controller, with user and group accounts stored outside the server. The server integrates with the domain controller infrastructure.

Note: The NAS server cannot act as a domain controller for other servers on the network. If user and group account information is stored locally, those accounts may be used only to authenticate logons to the NAS server, resulting in a workgroup configuration.

Administering users and groups in a domain environment is similar in a mechanical sense to administering them in a workgroup environment. If using an Active Directory domain controller, the Computer Management tool allows for adding, modifying, and removing users in the same context as in a workgroup environment. The concepts, however, are very different.

Additional information about planning for domain environments can be found at:

http://www.microsoft.com/windowsserver2003/technologies/directory/activedirectory/default.mspx

The configuration of the domain controller is reflected on the NAS server because it obtains user account information from the domain controller when deployed in a domain environment. As mentioned previously, the server cannot act as a domain controller itself.

User and Group Name Planning

Effective user and group management is dependent upon how well the user and group names are organized. Administrators typically create a small number of groups on the network and then assign users to the appropriate group or groups. File system and share permissions can then be applied at the group level, rather than at the user level. If the number of groups is small, assigning the appropriate permissions to selected group, or groups, is more efficient than assigning permissions to each user.

Although each organization has specific conventions, following general guidelines makes administration simpler and more efficient. Because CIFS/SMB is dependent on users and groups to grant appropriate access levels to file shares, CIFS/SMB administration benefits from a consistent user and group administration strategy.

Managing User Names

Usernames should reflect a logical relationship between the username and the person who uses the account. It is important that rules are established to ensure that usernames are:

- Systematic
- Easy to follow and implement
- Easy to remember

Using a combination of the user's first name, middle initial, and last name results in systematic usernames for every member of a particular organization. Common examples include:

- First initial followed by last name (jdoe for John Doe)
- First initial followed by middle initial and last name (jqpublic for John Q. Public)
- First name followed by last name, separated by a period (john.smith for John Smith)
- Last name followed by first initial (doej for Jane Doe)

Guidelines must be in place for instances when two users have the same initials or name. For example, a number can be added to the end of the username (jdoe1 and jdoe2).

Other conventions can be applied. Just ensure that conventions are both systematic and consistent.

Managing Group Names

Group management follows many of the same principles as user management.

It is recommended that group naming conventions be systematic and easy to understand. Make the group name convey some logical information about the function or purpose of the group. Table 11 provides examples of group names.

Table 11: Group Name Examples

Group Name	Description
Administrators	All designated administrators on the server
Users	All standard server users
Power users	All standard server users requiring advanced access levels

Using tags is a helpful convention that indicates the specific access that a particular user has to a network resource. For example, if there is a data share on the device, the network administrator can create a “Data Users ROnly” group and a “Data Users RWrite” group to contain users that have read only or read write access on the share, respectively.

Workgroup User and Group Management

Note: In a clustered environment, users and groups should not be managed locally.

In a workgroup environment, users and groups are managed through the WebUI of the NAS server. Within the Users option, there are two choices:

- Managing local users
- Managing local groups

User and group administrative tasks include adding, deleting, and modifying user and group information. Managing local users and managing local groups are discussed in the following paragraphs.

Managing Local Users

Managing users includes the following tasks:

- Adding a new user
- Deleting a user
- Setting a user password
- Modifying user properties

In the WebUI, under **Users**, **Local Users** is the **Local Users on Server** dialog box. All workgroup user administration tasks are performed in the **Local Users** dialog box.

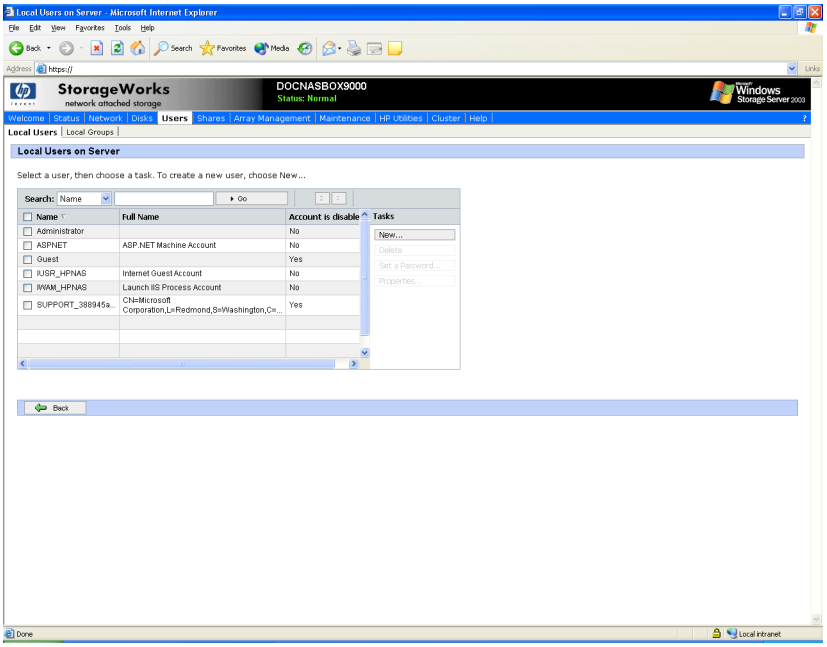

Figure 57: Local Users dialog box

All available options include: **New**, **Delete**, **Set a Password**, and **Properties**. When the **Local Users** dialog box is initially displayed, only the **New** option is available. After an existing user is selected, the additional actions are displayed. Each of these options is discussed in the following paragraphs.

Existing user records can be retrieved in one of two ways:

- By entering the user's User Name or Full Name in the Search fields to retrieve a specific user record. To redisplay the complete user list, space out the Search field.
- By selecting the user from the list of displayed users in the dialog box. The sort order of the display is controlled by clicking the Name field heading. The names are displayed in alphanumeric order or reverse alphanumeric order.

Adding a New User

To add a user:

1. From the **Local Users** dialog box, click **New**. The **Create New User** dialog box is displayed.

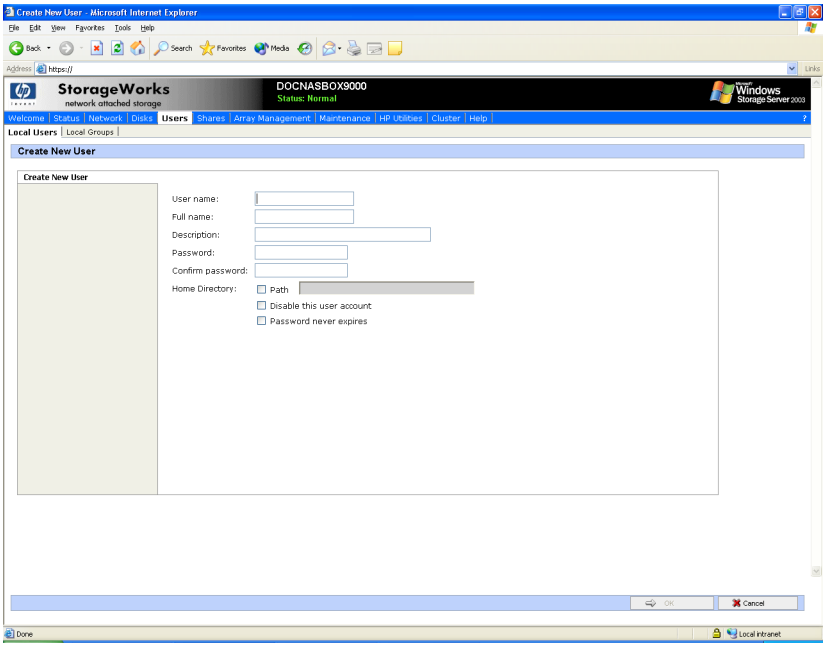

Figure 58: Create New User dialog box

2. Enter the user information and then click **OK**. The user is added and the **Local Users** dialog box is displayed again.

Deleting a User

To delete a user:

1. In the **Local Users** dialog box, select the user to delete, and then click **Delete**.

   The **Delete User** dialog box is displayed, including a warning note about deleting users.

2. To delete the user, click **OK**. The user is deleted and the **Local Users** dialog box is displayed again.

Modifying a User Password

Follow these steps to modify a user password:

1. In the **Local Users** dialog box, select the user whose password needs to be changed. Then, click **Set a Password**.

   The **Set Password** dialog box is displayed.

2. Enter the password and click **OK**. The Local Users dialog box is displayed again.

Modifying User Properties

To modify other user properties:

1. From the **Local Users** dialog box, select the user whose record needs to be modified. Then, click **Properties**.

   The General information page of the **Properties** dialog box is displayed. Figure 59 is an illustration of the **User Properties** dialog box.

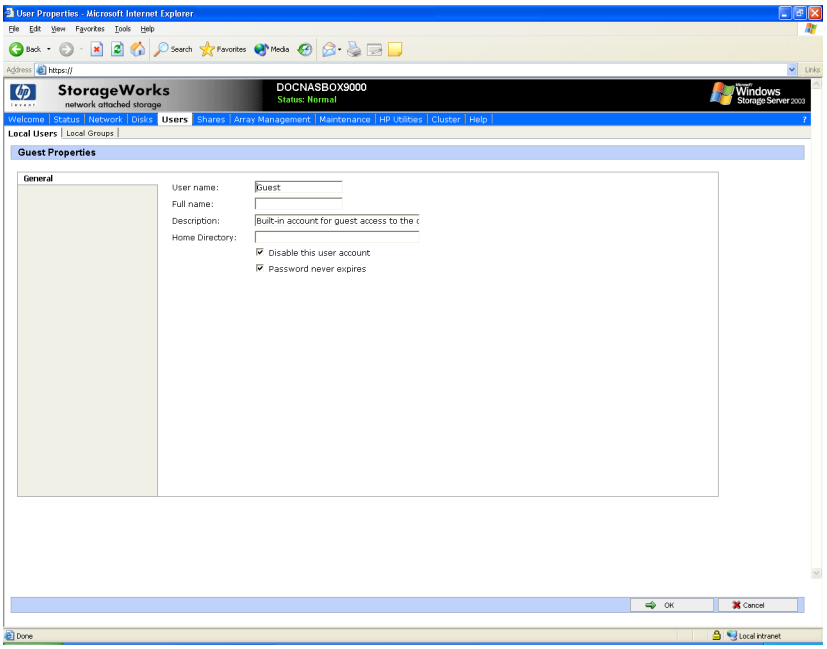

Figure 59: User Properties dialog box

2. The following information can be changed or set:
   - User name
   - Full name
   - Description
   - Home Directory
   - Disable this user account
   - Password expiration
3. After completing the changes, click **OK**. The **Local Users** dialog box is displayed again.

Managing Local Groups

Managing groups includes the following tasks:

- Adding a new group
- Deleting a group
- Modifying group properties, including user memberships

Local groups in a workgroup environment are managed through the Users option in the WebUI.

In the WebUI, under **Users**, **Local Groups** is the **Local Groups on Server** dialog box. All workgroup group administration tasks are performed in the **Local Groups on Server Appliance** dialog box.

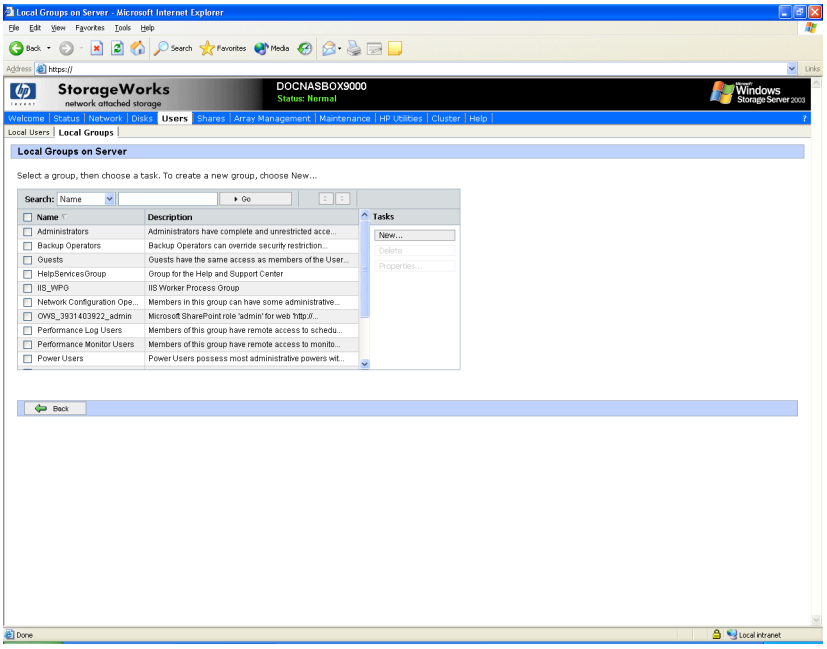

Figure 60: Local Groups dialog box

Adding a New Group

To add a group:

1. In the **Local Groups** dialog box, click **New**.

   The **Create New Group** dialog box is displayed.

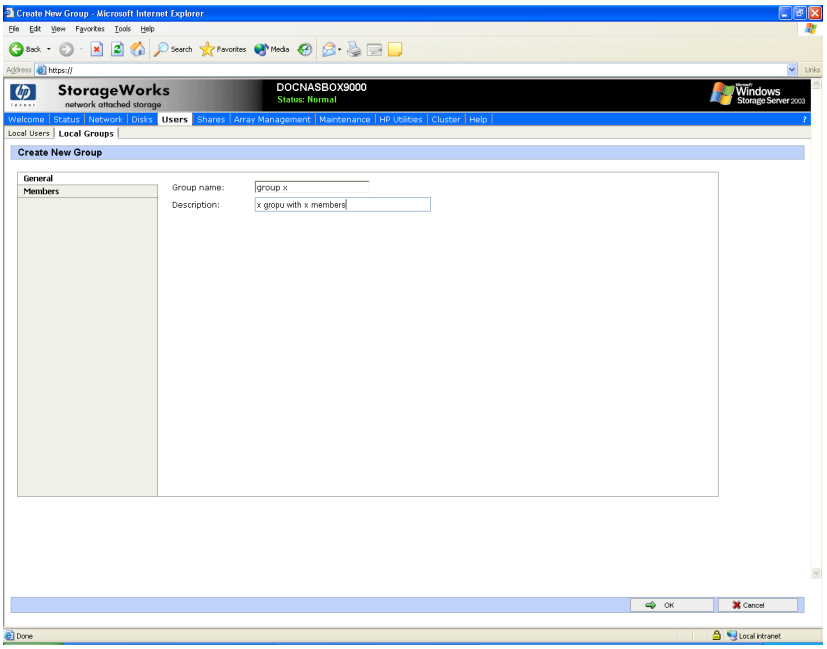

Figure 61: Create New Group dialog box, General tab

2. Enter the group name and description.
3. To indicate the user members of this group, click **Members**. See "Modifying Group Properties" for procedural instructions on entering group members.
4. After all group information is entered, click **OK**. The group is added, and the **Local Groups** dialog box is displayed again.

Deleting a Group

To delete a group:

1. From the **Local Groups** dialog box, select the group to delete, and then click **Delete**.
2. The **Delete Group** dialog box is displayed. Verify that this is the intended group and then click **OK**. The **Local Groups** dialog box is displayed again.

Modifying Group Properties

To modify other group properties:

1. From the **Local Groups** dialog box, select the desired group and then click **Properties**. The **Properties** dialog box is displayed.

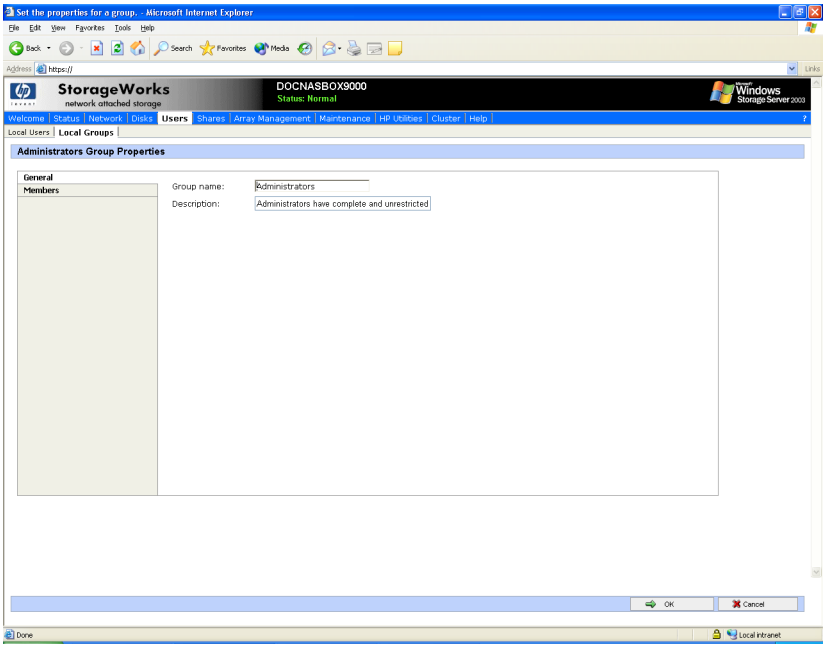

Figure 62: Group Properties dialog box, General tab

Within the Properties dialog box are two tabs:

- General tab
- Members tab

Each of these tabs is discussed in the following paragraphs.

2. Enter the desired changes in each of the tabs. Then, click **OK**. The **Local Groups** dialog box is displayed again.

General Tab

Within the General tab, basic group information can be changed, including:

- Group name
- Description

Members Tab

To indicate or change the members of a group, click the **Members** tab. Within this dialog box, users are added and removed from a group.

Two boxes are displayed: **Members** and **Add user or group**. Current members of that group are listed in the **Members** box. All users are listed in the **Add user or group** box.

- To add an existing local user to a group:
  1. Select the desired user from the **Add user or group** box
  2. Click the **Add** button.
  3. Click **OK** to save the changes.
- To remove an existing local user from a group:
  1. Select the desired user from the **Members** box.
  2. Click **Remove**.
  3. Click **OK** to save the changes.
- To add a user or group from a domain to this group, the scroll bar at the right of the screen may need to be used to scroll up the screen display:
  1. Enter the user or group name to include in the indicated format (domain\username).
  2. Select **Add**.
  3. Enter a domain\username and password.
  4. Click **OK** to complete adding the domain user or group.

Note: To add domain users and groups to a local group, the NAS device must be a member of the domain.

Figure 63 is an example of the **Members** tab.

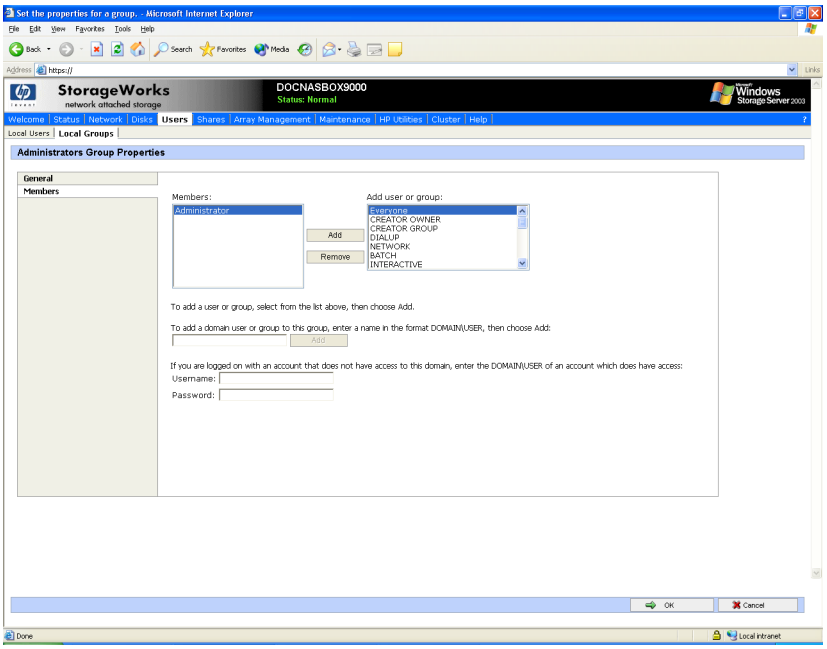

Figure 63: Group Properties dialog box, Members tab

Folder, Printer, and Share Management

7

The HP StorageWorks NAS server supports several file sharing protocols, including DFS, NFS, FTP, HTTP, and Microsoft SMB. This chapter discusses overview information as well as procedural instructions for the setup and management of the file shares for the supported protocols. In addition, discussions on security at the file level and at the share level are included in this chapter. As a new feature to the NAS server, printer services for network printers are now supported on the platform.

Abbreviated information on creating NFS file shares is included in this chapter; for detailed information on setting up and managing NFS file shares, see the "Microsoft Services for NFS" chapter.

NCP shares must be set up and managed through the NAS Management Console user interface. For information on managing NCP file shares, see the "NetWare File System Management" chapter.

More information about Windows file system security is available on the Microsoft website:

www.microsoft.com

Note: The NAS 4000s and 9000s servers can be deployed in a clustered as well as a non-clustered configuration. This chapter discusses share setup for a non-clustered deployment. For information on managing file shares and printers in a cluster, see the Cluster Administration chapter.

All procedures in this chapter are documented using the WebUI. In addition to this guide, you may use the WebUI online help.

Folder Management

Volumes and folders on any system are used to organize data. Regardless of system size, systematic structuring and naming conventions of volumes and folders eases the administrative burden. Moving from volumes to folders to shares increases the level of granularity of the types of data stored in the unit and the level of security access allowed.

Although a variety of methods can be used to create and manage file folders on the NAS server, this document discusses using the NAS Web based user interface (WebUI.)

Managing system volumes and file folders includes the following tasks:

- Navigating to a specific volume or folder
- Creating a new folder
- Deleting a folder

- Modifying folder properties
- Creating a new share for a volume or folder
- Managing shares for a volume or folder
- Managing file level permissions

Navigating to a Specific Volume or Folder

When you work with volumes and folders, the first task is to gain access to the desired volume or folder.

The steps are the same, whether navigating to a volume or a folder:

1. To navigate to a specific volume or folder, from the WebUI, select **Shares** and then **Folders**. Initially, the **Volumes** dialog box is displayed.

   This initial dialog box displays all system volumes.

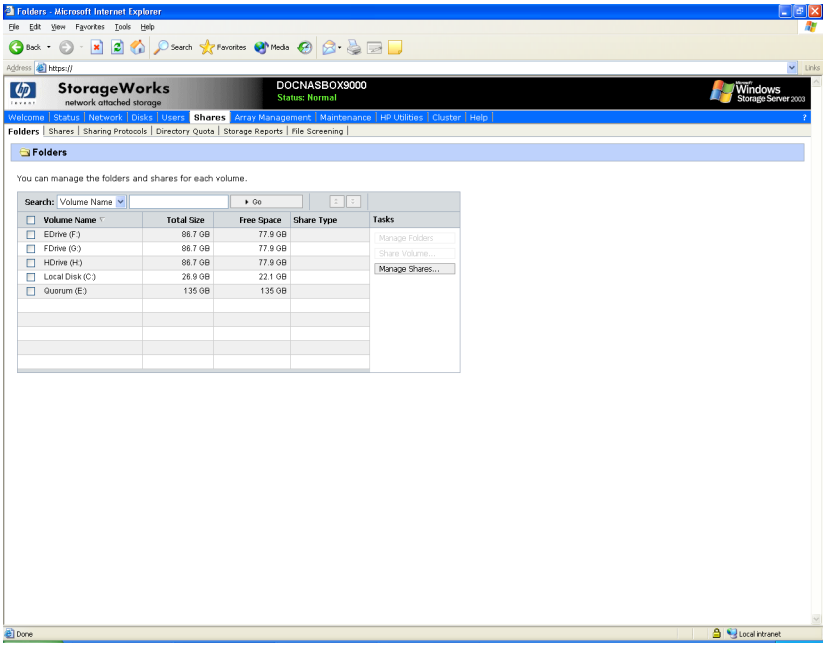

Figure 64: Volumes dialog box

2. From this dialog box, navigate to a specific folder by selecting the appropriate volume and then clicking **Manage Folders**. The **Folders** dialog box is displayed, with a list of all of the folders within that volume.
3. To navigate to a subfolder, select the folder in which the subfolder resides, and then click **Open**. Repeat this searching and opening process until the desired folder is opened. See Figure 65 for an example of **Folders** dialog box.

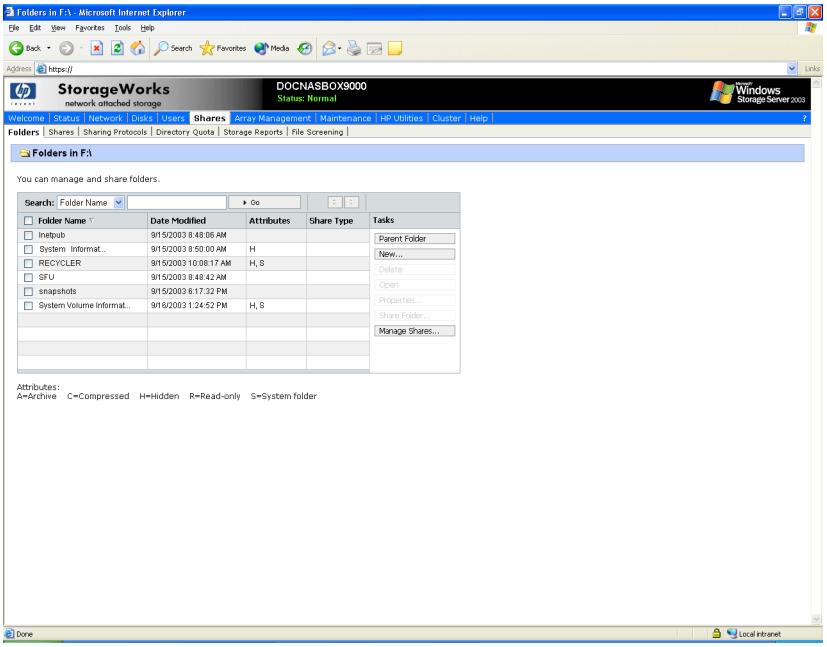

Figure 65: Folders dialog box

After accessing the desired folder, the following actions can be performed:

- Creating a new folder
- Deleting a folder
- Modifying folder properties
- Creating a new share for the volume or folder
- Managing shares for the volume or folder

Creating a New Folder

To create a new folder:

1. From the **Shares** directory, navigate to the **Folders** menu and then select **New**. The **Create New Folder** dialog box is displayed.

   Two tabs are displayed: **General** and **Compress**. Use these two tabs to enter the parameters for the new folder.

2. In the General tab, enter a name for the folder and specify the folder attributes.

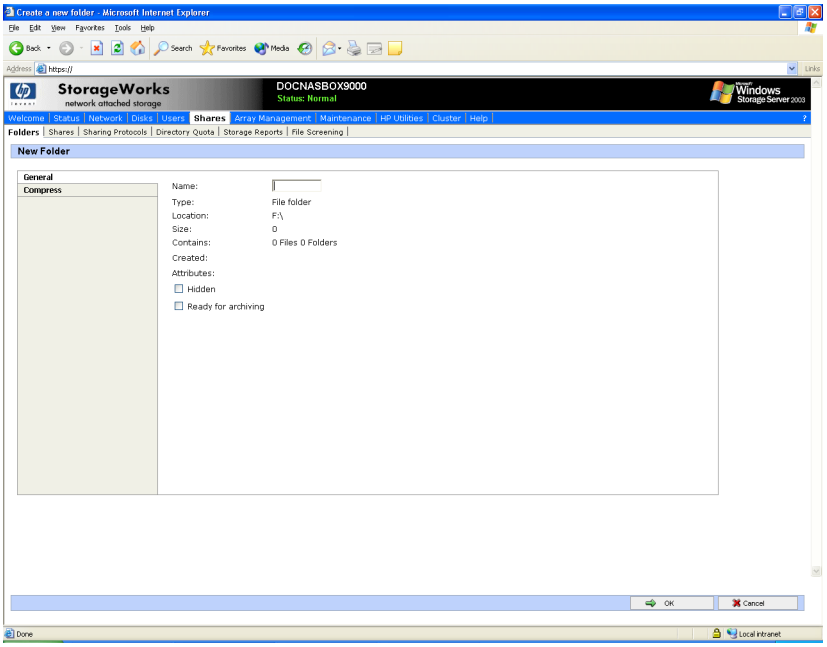

Figure 66: Create a New Folder dialog box, General tab

3. In the **Compress** tab, indicate whether and how this folder and its contents are to be compressed.
4. After all information for the new folder is entered, click **OK**.

Deleting a Folder

To delete a folder:

1. From the **Shares** directory, navigate to the folder to delete. Select the folder and then click **Delete**. The **Delete Folder** dialog box is displayed.

   Summary information about the deletion is displayed.

Note: View the summary information to confirm that this is the intended share.

---

2. Verify that the displayed folder is the folder to delete and then click **OK**.

   The folder and all of its subfolders are deleted and the main dialog box is displayed again.

Modifying Folder Properties

To modify folder properties:

1. From the **Shares** directory, navigate to the folder whose properties need to be edited. Then click **Properties**. The **Properties** dialog box is displayed.

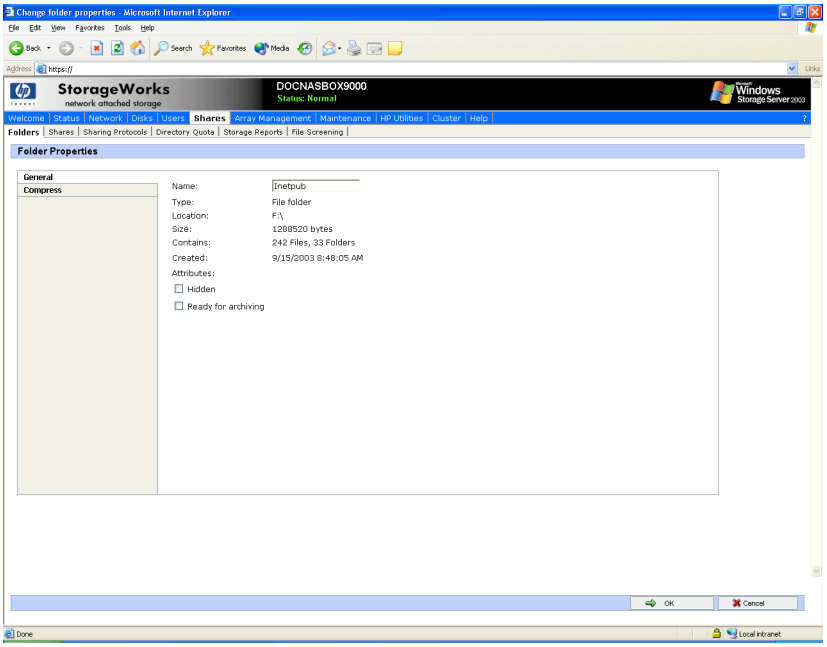

Figure 67: Folder Properties dialog box, General tab

2. In the **General** tab, enter the new information for the folder, which may include:
   - Folder Name
   - Folder Attributes
3. In the **Compress** tab, indicate whether and how this folder and its contents are to be compressed.
4. After all changes have been completed, click **OK**. The **Folders** dialog box is displayed again.

Creating a New Share for a Volume or Folder

Within the WebUI, there are two access points to the same screens used to create file shares:

- A share can be created for a folder while working with that folder in the **Folders** screens.
- A share can be created and, if necessary, new folders can be created, while working with file shares in the **Shares** screens.

This section discusses creating shares from the **Folders** menu, and is an overview of the procedures. Complete details on the process of creating shares are included in the discussion that documents creating shares through the **Shares** menu. See the “Managing Shares” section of this chapter for these details.

Note: This function will operate in a cluster but should only be used for non-cluster aware shares. Use Cluster Administrator to create shares for a cluster.

To create a new share for a specific volume or folder while in the **Folders** menu:

1. Navigate to the desired volume or folder and click **Manage Shares**.

2. Click **New**. The **Create New Share** dialog box is displayed.

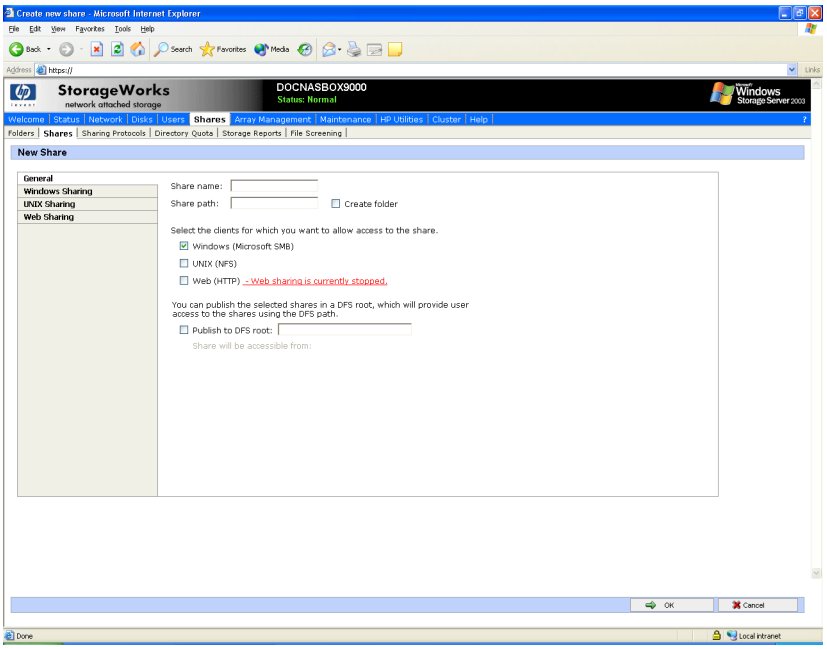

Figure 68: Create New Share dialog box, General tab

3. Enter the information for the share, including the name of the share, the allowed protocols, and corresponding permissions.

Note: The **Share path** is the path of the previously selected volume or folder. This field is automatically completed by the system.

---

4. Select the appropriate tab to enter protocol specific information.

   See the "Managing Shares" section for detailed information about these entries.

5. After entering all share information, click **OK**.

Note: The default permission settings for a new share are read-only.

Managing Shares for a Volume or Folder

Within the WebUI, there are two access points to the same screens used to manage file shares:

- While working with a folder in the **Folders** dialog boxes, the administrator can create, delete, and modify shares for that folder.
- While working with file shares in the **Shares** dialog boxes, the administrator can create, delete, and modify shares (and if necessary, create new folders).

Note: This section discusses managing shares from the **Folders** menu, and is an overview of the procedures. Complete details on the process of managing shares are included in the discussion that documents creating shares through the **Shares** menu. See the "Managing Shares" section later in this chapter for these details.

To create, delete, and manage shares for a particular volume or folder while in the **Folders** menu:

1. From the **Folders** directory, navigate to the target volume or folder and click **Manage Shares**. The **Shared Folders** dialog box is displayed.

   All associated shares for that folder or volume are listed.

2. To create a new share, click **New**. The **Create a New Share** dialog box is displayed.

   Because the screens are the same whether shares are managed through the **Folders** menu or the **Shares** menu, the procedures are only documented once. See "Creating a New Share" in the "Share Management" section for detailed procedural instructions on creating new file shares.

3. To delete a share, select the share to delete and click **Delete**. The **Delete Share** dialog box is displayed.

   Because the screens are the same whether shares are managed through the **Folders** menu or the **Shares** menu, the procedures are only documented once. See "Deleting a New Share" in the "Share Management" section for detailed procedural instructions on deleting file shares

4. To modify share properties, select the share to modify, and click **Properties**. The **Share Properties** dialog box is displayed.

   Because the screens are the same whether shares are managed through the **Folders** menu or the **Shares** menu, the procedures are only documented once. See "Modifying Share Properties" in the "Share Management" section for detailed procedural instructions on modifying shares.

Managing File Level Permissions

The WebUI of the NAS server provides security at the share level and is discussed later in this chapter. Security at the file level is managed using Windows Explorer available from the Desktop of the NAS server. To access the NAS server Desktop from the WebUI, go to the **Maintenance** menu and select **Remote Desktop**.

File level security includes settings for permissions, ownership, and auditing for individual files.

To enter file permissions:

1. Using Windows Explorer, navigate to the folder or file that needs to be changed and then right-click the folder.
2. Select **Properties**, select the **Security** tab, then click **Advanced**. Figure 69 illustrates the properties available on the **Advanced Security Settings** dialog box.

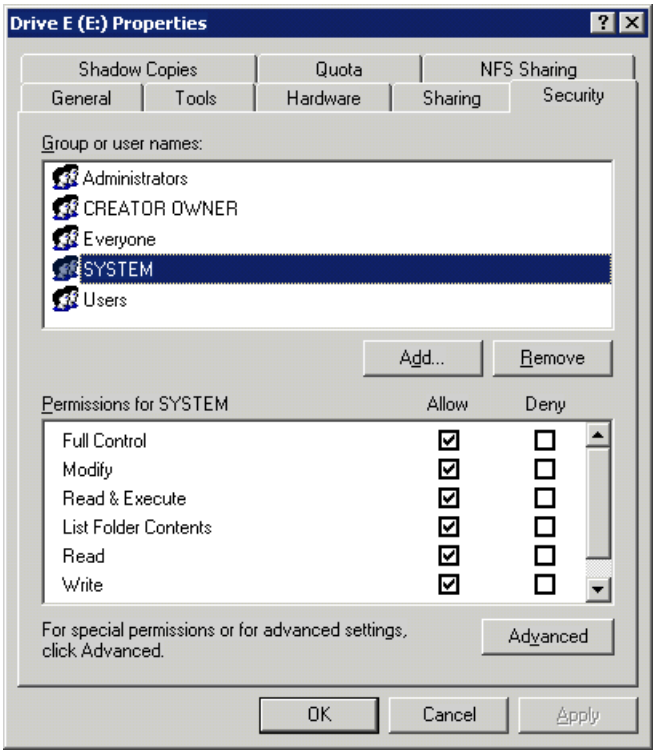

Figure 69: Security Properties dialog box

Several options are available in the **Security** tab dialog box:

- To add users and groups to the permissions list, click **Add**. Then follow the dialog box instructions.
- To remove users and groups from the permissions list, highlight the desired user or group and then click **Remove**.
- The center section of the **Security** tab provides a listing of permission levels. When new users or groups are added to the permissions list, select the appropriate boxes to configure the common file access levels.
- To modify ownership of files or to modify individual file access level permissions, click **Advanced**.

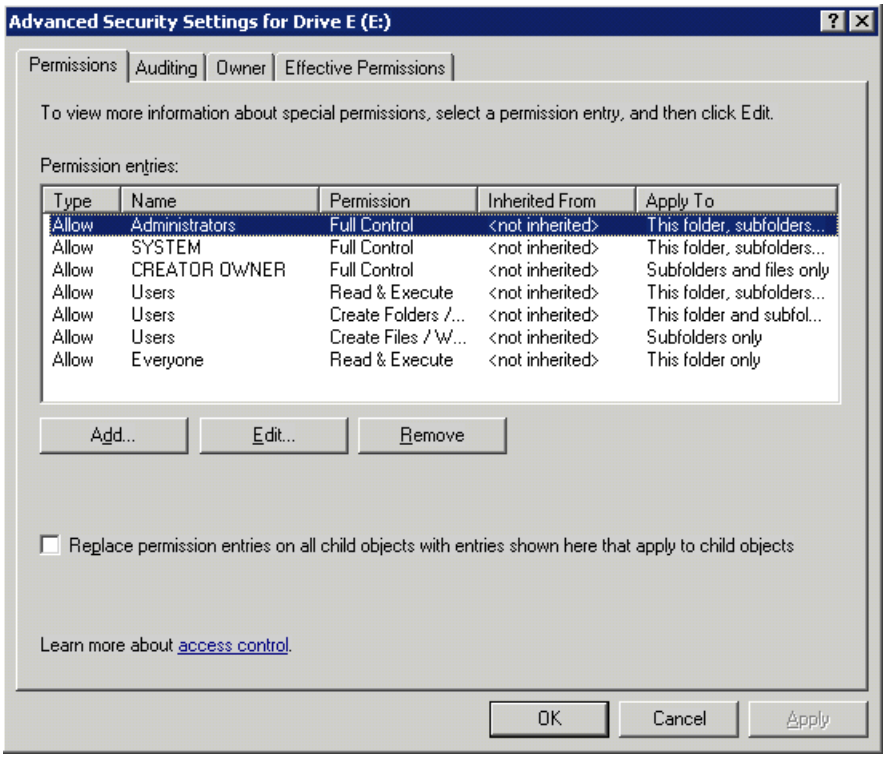

Figure 70: Advanced security settings

To modify specific permissions assigned to a particular user or group for a selected file or folder in the **Advanced** screen:

1. Select the desired user or group.
2. Click **Edit**.
3. Check all the permissions that you want to enable, and clear the permissions that you want to disable. Enable or disable permissions by selecting the **Allow** box to enable permission or the **Deny** box to disable permission. If neither box is selected, permission is automatically disabled. Figure 71 illustrates the **Edit** screen and some of the permissions.

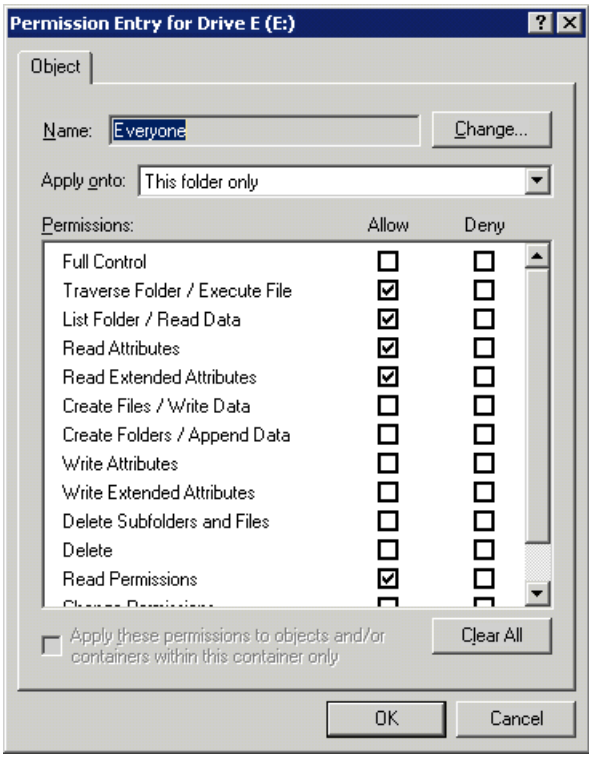

Figure 71: User or Group Permission Entry dialog box

Other functionality available in the **Advanced Security Settings** tab is illustrated in Figure 70 and includes:

- **Add a new user or group**. Click **Add**, and then follow the dialog box instructions.
- **Remove a user or group**. Click **Remove**.
- **Replace permission entries on all child objects with entries shown here that apply to child objects**. This allows all child folders and files to inherit the current folder permissions by default.

Another area of the **Advanced Security Settings** is the **Auditing** tab. Auditing allows you to set rules for the auditing of access, or attempted access, to files or folders. Users or groups can be added, deleted, viewed, or modified through the advanced **Advanced Security Settings Auditing** tab. The **Auditing** tab dialog box is illustrated in Figure 72.

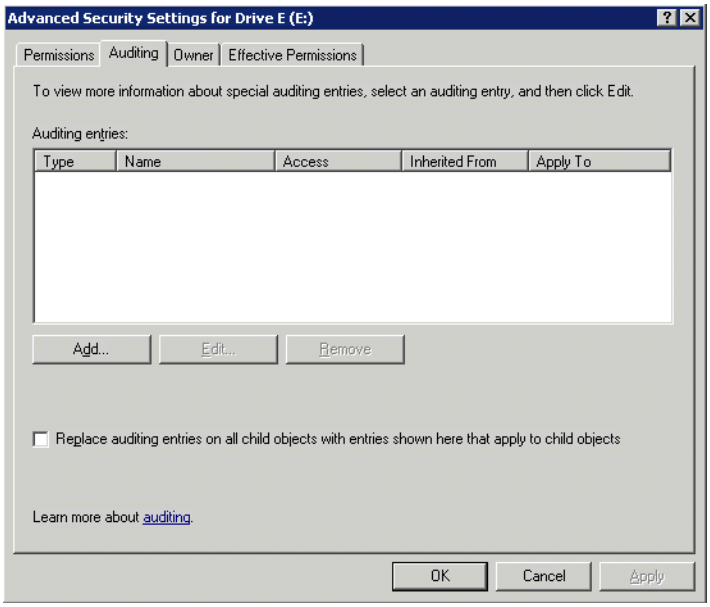

Figure 72: Advanced Security Settings, Auditing tab dialog box

4. Click **Add** to display the Select User or Group dialog box.

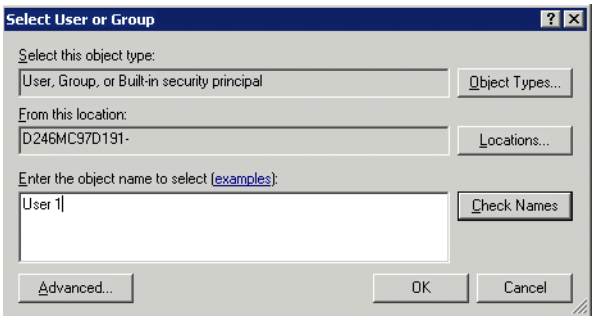

Figure 73: Select User or Group dialog box

Note: Click Advanced to search for users or groups.

5. Select the user or group.

6. Click **OK**. Figure 74 illustrates the **Auditing Entry** screen that is displayed.

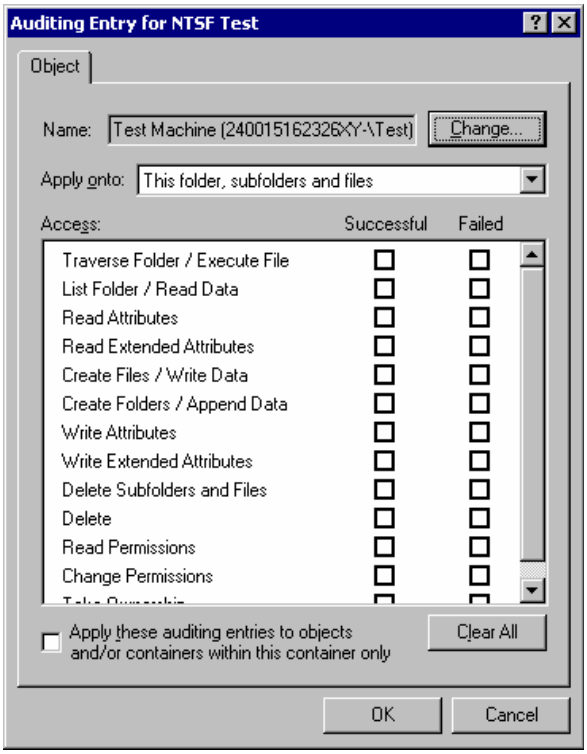

Figure 74: Auditing Entry dialog box for folder name NTSF Test

7. Select the desired **Successful** and **Failed** audits for the user or group as shown in Figure 74.
8. Click **OK**.

Note: Auditing must be enabled to configure this information. Use the local Computer Policy Editor to configure the audit policy on the NAS server.

The **Owner** tab allows for taking ownership of files. Typically, administrators use this area to take ownership of files when the file ACL is incomplete or corrupt. By taking ownership, you gain access to the files and then manually apply the appropriate security configurations. Figure 75 illustrates the **Owner** tab.

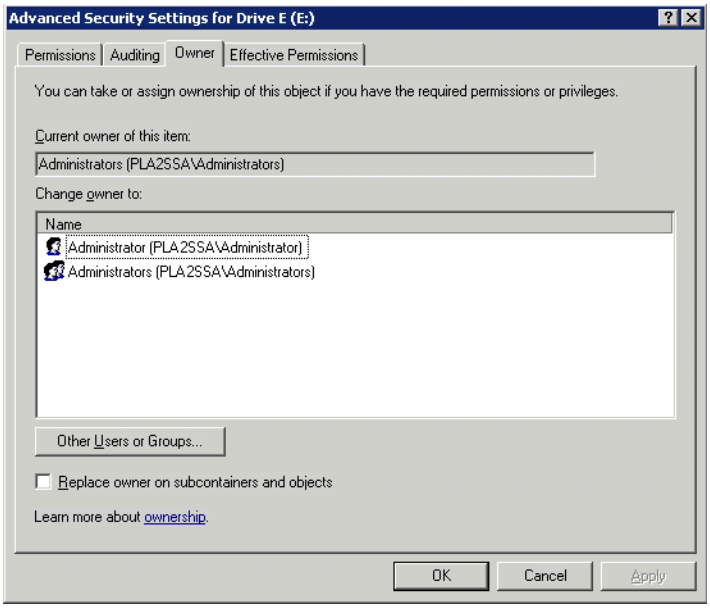

Figure 75: Advanced Security Settings, Owner tab dialog box

The current owner of the file or folder is listed at the top of the screen. To take ownership:

1. Select the appropriate user or group from the **Change owner to** list.
2. If it is also necessary to take ownership of subfolders and files, enable the **Replace owner on subcontainers and objects** box.
3. Click **OK** to execute the commands.

Share Management

There are several ways to set up and manage shares. The WebUI provides screens for setting up and managing shares. Additional methods include using a command line interface, Windows Explorer, or NAS Management Console. This guide demonstrates using the WebUI to set up and manage shares.

Note: The NAS 4000s and 9000s servers can be deployed in a clustered as well as a non-clustered configuration. This chapter discusses share setup for a non-clustered deployment. For information on managing file shares and printers in a cluster, see the Cluster Administration chapter.

As previously mentioned, the file sharing security model of the NAS device is based on the NTFS file-level security model. Share security seamlessly integrates with file security. In addition to discussing share management, this section discusses share security. See "Managing File Level Permissions" earlier in this chapter for information on file security.

Shares management topics include:

- Share Considerations
- Defining Access Control Lists
- Integrating Local File System Security into Windows Domain Environments
- Comparing Administrative and Standard Shares
- Planning for Compatibility between File-Sharing Protocols
- Managing Shares

Share Considerations

Planning the content, size, and distribution of shares on the NAS server can improve performance, manageability, and ease of use.

The content of shares should be carefully chosen to avoid two common pitfalls: either having too many shares of a very specific nature or of having very few shares of a generic nature. For example, shares for general usage are easier to set up in the beginning, but can cause problems later. Frequently, a better approach is to create separate shares with a specific purpose or group of users in mind. However, creating too many shares also has its drawbacks. Take care to avoid creating shares unnecessarily. For example, if it is sufficient to create a single share for user home directories, create a "homes" share rather than creating separate shares for each user.

By keeping the number of shares and other resources low, the performance of the NAS server is optimized. For example, instead of sharing out each individual user's home directory as its own share, share out the top level directory and let the users map personal drives to their own subdirectory.

Defining Access Control Lists

The Access Control List (ACL) contains the information that dictates which users and groups have access to a share, as well as the type of access that is permitted. Each share on an NTFS file system has one ACL with multiple associated user permissions. For example, an ACL can define that User1 has read and write access to a share, User2 has read only access, and User3 has no access to the share. The ACL also includes group access information that applies to every user in a configured group. ACLs are also referred to as permissions.

Integrating Local File System Security into Windows Domain Environments

ACLs include properties specific to users and groups from a particular workgroup server or domain environment. In a multidomain environment, user and group permissions from several domains can apply to files stored on the same device. Users and groups local to the NAS server can be given access permissions to shares managed by the device. The domain name of the NAS server supplies the context in which the user or group is understood. Permission configuration depends on the network and domain infrastructure where the server resides.

File-sharing protocols (except NFS) supply a user and group context for all connections over the network. (NFS supplies a machine based context.) When new files are created by those users or machines, the appropriate ACLs are applied.

Configuration tools provide the ability to share permissions out to clients. These shared permissions are propagated into a file system ACL and when new files are created over the network, the user creating the file becomes the file owner. In cases where a specific subdirectory of a share has different permissions from the share itself, the NTFS permissions on the subdirectory apply instead. This method results in a hierarchical security model where the network protocol permissions and the file permissions work together to provide appropriate security for shares on the device.

Note: Share permissions and file level permissions are implemented separately. It is possible for files on a file system to have different permissions from those applied to a share. When this situation occurs, the file level permissions override the share permissions.

Comparing Administrative (Hidden) and Standard Shares

CIFS supports both administrative shares and standard shares. Administrative shares are shares with a last character of $. Administrative shares are not included in the list of shares when a client browses for available shares on a CIFS server. Standard shares are shares that do not end in a $ character. Standard shares are listed whenever a CIFS client browses for available shares on a CIFS server.

The NAS server supports both administrative and standard CIFS shares. To create an administrative share, end the share name with the $ character when setting up the share. Do not type a $ character at the end of the share name when creating a standard share.

Planning for Compatibility between File Sharing Protocols

When planning for cross-platform share management on the NAS server, it is important to understand the different protocols and their associated constraints. Each additional protocol that is supported adds another level of constraints and complexity.

NFS Compatibility Issues

When planning to manage CIFS and NFS shares, consider two specific requirements.

Note: Further information, including details about the NFS Service and the User Mapping service, is available in the "Microsoft Services for NFS" chapter.

- **NFS service does not support spaces in the names for NFS file share**s.

  NFS translates any spaces in an export into an underscore character. Additional translations can be set up for files. See the "OEM Supplemental Help" chapter of the SFU help, found on the NAS server. This feature is designed to ensure the greatest level of compatibility with NFS clients, because some do not work with NFS exports that contain a space in the export name.

  If you plan to use the same name when sharing a folder through CIFS, and then exporting it through NFS, do not put spaces in the CIFS share name.

- **NFS service does not support exporting a child folder when its parent folder has already been exported**.

  An NFS client can access a child folder by selecting the parent folder and then navigating to the child folder. If strict cross-platform compatibility is an administration goal, CIFS must be managed in the same way. Do not share a folder through CIFS if the parent folder is already shared.

Managing Shares

Shares can be managed through the **Shares** menu option of the WebUI. Tasks include:

- Creating a new share
- Deleting a share
- Modifying share properties

Each of these tasks is discussed in this section.

Note: These functions will operate in a cluster but should only be used for non-cluster aware shares. Use Cluster Administrator to manage shares for a cluster. The page will display cluster share resources.

Creating a New Share

To create a new share:

1. From WebUI main menu, select the **Shares** directory and then select the **Shares** option. The **Shares** dialog box is displayed. From the **Shares** dialog box, click **New**. The **General** tab of the **Create a New Share** dialog box is displayed.

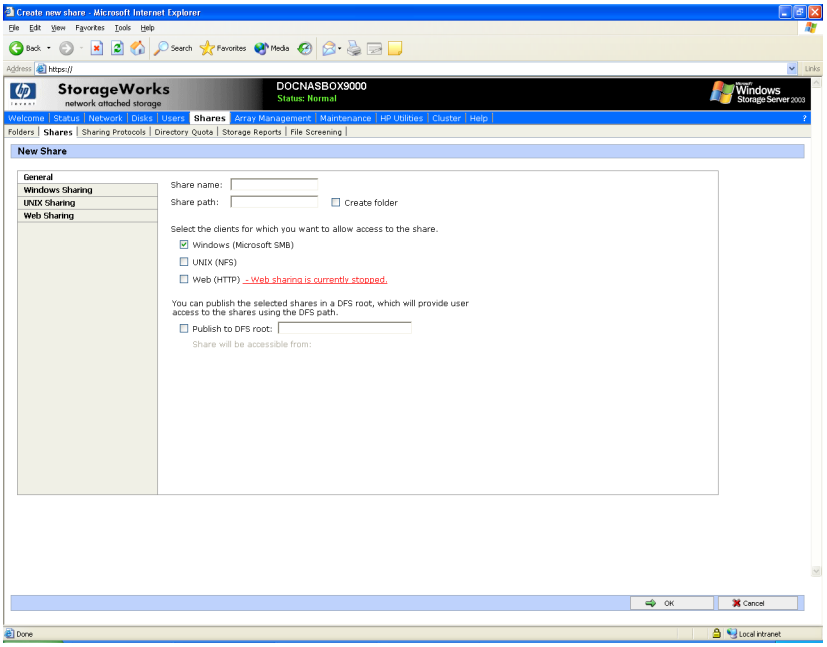

Figure 76: Create a New Share dialog box, General tab

2. Enter the following information:
   - Share name
   - Share path
   - Client protocol types

   To create a folder for the new share, check the indicated box and the system will create the folder at the same time it creates the share.

   Protocol specific tabs are available to enter sharing and permissions information for each sharing type. See “Modifying Share Properties” for detailed information on these tabs.
3. After entering all share information, click **OK**.

Deleting a Share

Caution: Before deleting a share, warn all users to exit that share and confirm that no one is using the share.

To delete a share:

1. From the **Shares** menu, click **Delete**.
2. Verify that this is the correct share, and then click **OK**.

Modifying Share Properties

To change share settings:

1. From the **Shares** menu, select the share to modify and then click **Properties**. The **General** tab of the **Share Properties** dialog box is displayed.

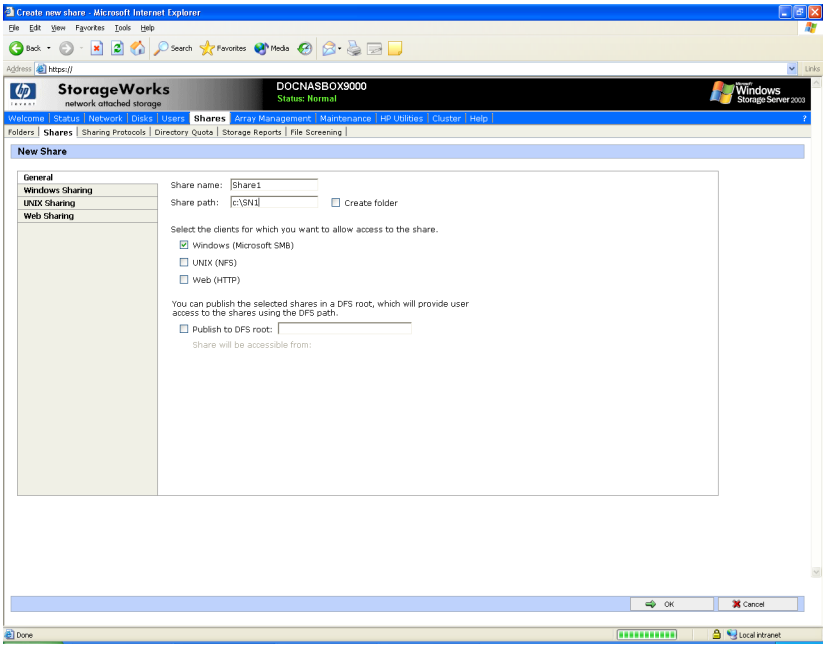

Figure 77: Share Properties dialog box, General tab

The name and path of the selected share is displayed.

2. To enter or change client protocol information, check the appropriate boxes and then click the corresponding tabs.
   - Windows Sharing
   - UNIX Sharing
   - Web Sharing (HTTP)

   Each of these tabs is discussed in the following paragraphs.
3. After all share information has been entered, click **OK**. The **Share** menu is displayed again.

Windows Sharing

From the **Windows Sharing** tab of the **Share Properties** dialog box:

1. Enter a descriptive **Comment**, and the **User limit** (optional).

   See Figure 78 for an example of the **Windows Sharing** tab screen display.

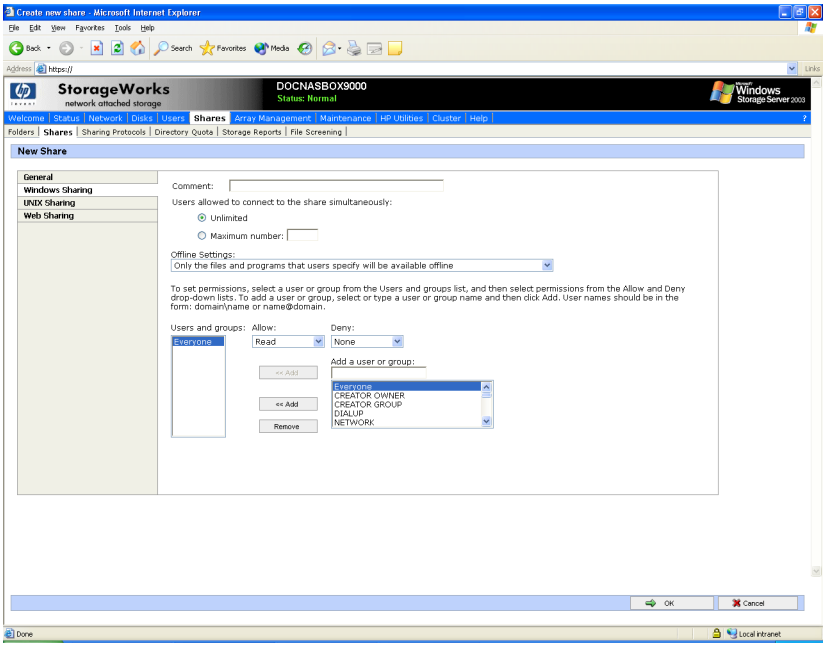

Figure 78: Share Properties dialog box, Windows Sharing tab

2. Select Offline settings.
3. Set the permissions.

   The **Permissions** box lists the currently approved users for this share.

   - To add a new user or group, either select a user or group from the box at the bottom right of the screen or manually enter the user or group name in the Add a user or group box and then click Add. That user or group is added to the Permissions box.
   - To remove access to a currently approved user or group, select the user or group from the Permissions box and then click Remove.
   - To indicate the type of access allowed for each user, select the user and then expand the Allow and Deny drop down boxes. Select the appropriate option.
4. After all Windows Sharing information is entered, either click the next **Sharing** tab or click **OK**.

UNIX Sharing

From the **UNIX Sharing** tab of the **Create a New Share** dialog box:

1. Indicate the machines that will have access to this share.

   Select the machine to include in the **Select a group** box or manually enter the NFS client computer name or IP address. Then click **Add**.

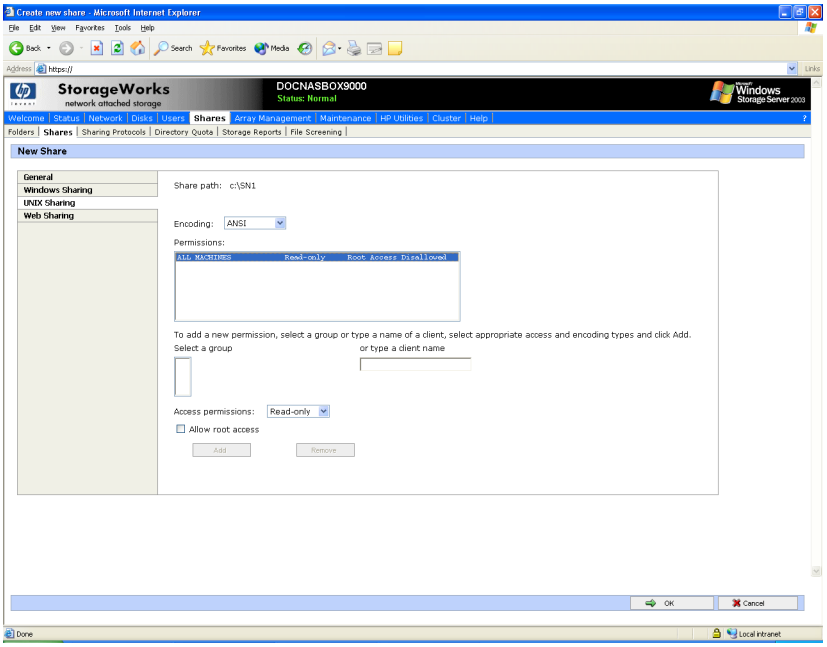

Figure 79: Share Properties dialog box, UNIX Sharing tab

2. Indicate the access permissions.

   Select the machine from the main user display box and then select the appropriate access method from the **Access permissions** drop down box.

   The types of access are:

   - **Read-only**—Use this permission to restrict write access to the share.
   - **Read-write**—Use this permission to allow clients to read or write to the share.
   - **No access**—Use this permission to restrict all access to the share.

3. Select whether or not to allow root access.

   - **Read-only + Root**—Use this permission to restrict write access to the share. Use this permission to assign administrative access to the share. This will allow the client computer to have root access to the NFS share. Map the UNIX root user to the Windows user Administrator. Also, map the group that this UNIX root belongs to, to the Windows group Administrator.
   - **Read-write + Root**—Use this permission to allow clients to read or write to the share. Use this permission to assign administrative access to the share. This will allow the client computer to have root access to the NFS share. Map the UNIX root user to the Windows user Administrator. Also, map the group that this UNIX root belongs to, to the Windows group Administrator.

4. After all UNIX sharing information is entered, click **OK**.

Web Sharing (HTTP)

From the **Web Sharing** tab of the **Create New Share** dialog box:

1. Select the read and write access permissions that are allowed.
2. Click **OK**.

Protocol Parameter Settings

As previously mentioned, the NAS server supports the following protocols:

- DFS
- NFS
- FTP
- HTTP
- Microsoft SMB

This section discusses the parameter settings for each protocol type.

Note: See the protocol section of the Cluster Administration chapter for information about protocol selection and management in a cluster.

To access and enter protocol parameter settings:

1. From the **Shares** menu, select **Sharing Protocols**. The **File-Sharing Protocols** dialog box is displayed.

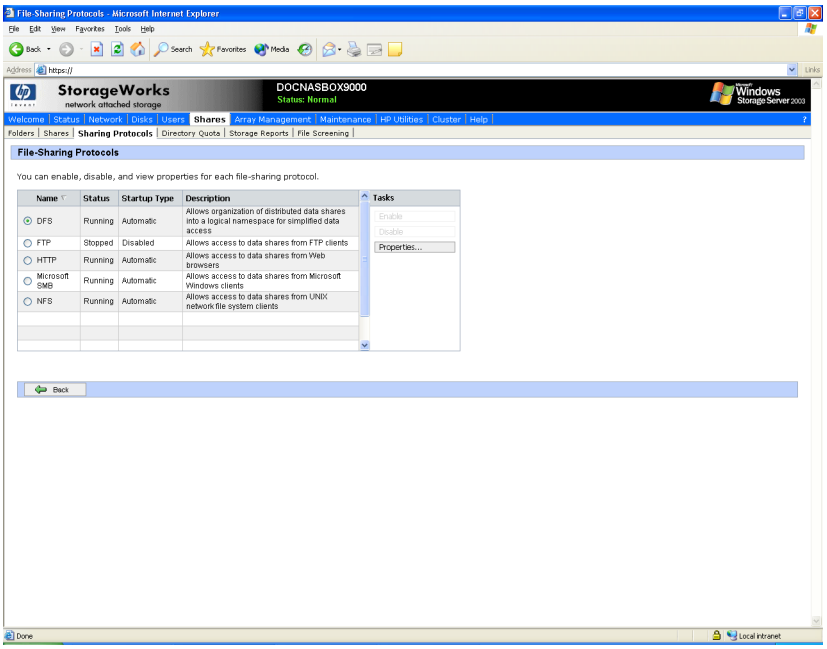

Figure 80: File-Sharing Protocols dialog box

2. Protocols and their statuses are listed. The following options are available:
   - Enabling a protocol
   - Disabling a protocol
   - Modifying Protocol Settings

Because enabling and disabling a protocol are self explanatory, only modifying protocol specific settings is described in this section.

DFS Protocol Settings

With Distributed File System (DFS) and the Windows SMB protocol, files can be distributed across multiple servers and appear to users as if they reside in one place on the network. A configuration containing multiple shares is known as a virtual namespace.

Using Distributed File System (DFS), system administrators can make it easy for users to access and manage files that are physically distributed across a network. Users do not need to know and specify the actual physical location of files in order to access them.

For example, if documents are scattered across multiple servers in a domain, DFS can make it appear as though the documents all resides on a single server. This eliminates the need for users to go to multiple locations on the network to find the information.

Each DFS namespace requires a root. A DFS root is a starting point of the DFS namespace. The root is often used to refer to the namespace as a whole. A root maps to one or more root targets, each of which corresponds to a shared folder on a server. A root is implemented as a shared folder on the DFS server.

Deploying DFS

A distributed file system can be implemented as a stand-alone root distributed file system or as a domain root distributed file system. The type of a distributed file system determines which client computers can access the distributed file system.

A stand-alone DFS root:

- Does not use Active Directory to manage DFS
- Cannot have more than one root on a server
- Does not support automatic file replication using the File Replication service (FRS)
- Is not fault tolerant and if the root fails the entire namespace will collapse.

A domain DFS root:

- Must be hosted on a domain member server
- Has its DFS namespace automatically published to Active Directory
- Can have more than one root on a server
- Supports automatic file replication through FRS
- Supports fault tolerance through FRS

Two points of management of the DFS namespace are provided with the NAS server. These points of management are the WebUI and the Distributed File System Administration Tool located on the local console of the NAS server under **Start > Programs > Administrative Tool**. See Figure 81. The WebUI is designed to provide the following functions:

- Stand alone root management (Add, Delete)
- Share publishing to stand alone or domain DFS
- Default behavior for DFS share publishing

All other functions must be performed via the DFS Administration Tool. The NAS server administration guide only provides instructions on the Web UI portion of the product. The DFS Administration Tool is complete with online help. In addition, general information on DFS may be found at:

http://www.microsoft.com/windowsserver2003/techinfo/overview/dfs.mspx

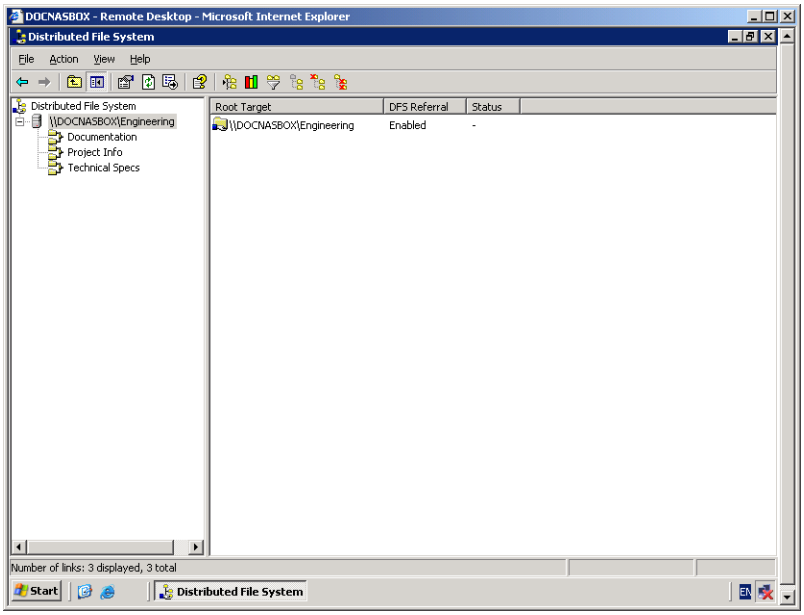

Figure 81: DFS Win32 GUI

DFS Administration Tool

The DFS Administration Tool provides extended functionality not found in the WebUI. These functions include:

- Management of multiple DFS Roots on multiple machines from a single interface
- Domain based DFS management
- Target and Link management
- Status Checks of a DFS managed share link
- Exporting of the DFS names space to a text file

The NAS server administration guide only provides instructions on the WebUI portion of the product. The DFS Administration Tool is complete with online help. In addition, general information on DFS may be found at:

http://www.microsoft.com/windowsserver2003/techinfo/overview/dfs.mspx

Accessing the DFS Namespace from other Computers

In addition to the server-based DFS component of the Windows Storage Server 2003 family, there is a client-based DFS component. The DFS client caches a referral to a DFS root or a DFS link for a specific length of time, defined by the administrator.

The DFS client component runs on a number of different Windows platforms. In the case of older versions of Windows, the client software must be downloaded to run on that version of Windows. Newer versions of Windows have client software built-in.

Non-Windows (such as Linux/UNIX) based clients can not access the DFS namespace as DFS is dependent on a Windows component to function.

Setting DFS Sharing Defaults

The We UI can be used to set the default DFS settings provided when creating a shared folder. When a new shared folder is create, the DFS defaults may be overridden.

To set DFS sharing defaults:

1. On the primary navigation bar, choose **Shares**.
2. Choose **Sharing Protocols**.
3. Select **DFS**, and then choose **Properties**.

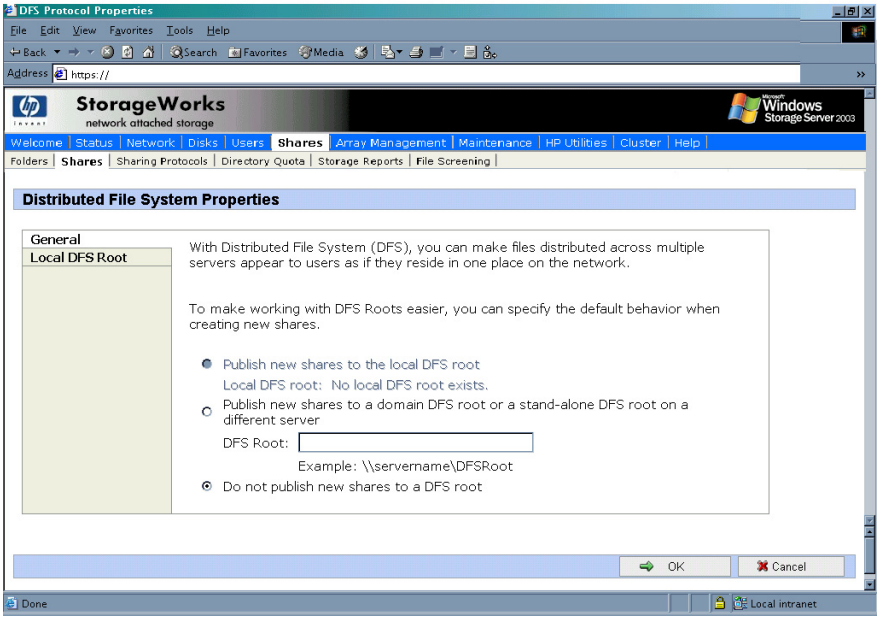

Figure 82: DFS properties, general tab

4. On the General tab, choose the default settings that are desired when creating a shared directory.
   - To set the default to publish the share to the local DFS root, select **Publish new shares to the local DFS root**.
   - To set the default to publish the share to another DFS root, select **Publish new shares to a domain DFS root or a stand-alone DFS root on a different server**. In the DFS root box, type the path of the default DFS root.
   - To not publish the share to a DFS root, select **Do not publish new shares to a DFS root**.
5. Choose **OK**.

Creating a Local DFS Root

The WebUI can be only be used to create a single, local stand-alone DFS root on the server as mentioned previous. To create a local domain DFS root use the DFS administrative tool. For more information about DFS root types refer to the section above entitled “Deploying DFS”.

To create a local stand-alone DFS root:

1. On the primary navigation bar, choose **Shares**.
2. Choose **Sharing Protocols**.

3. Select **DFS**, and then choose **Properties**.

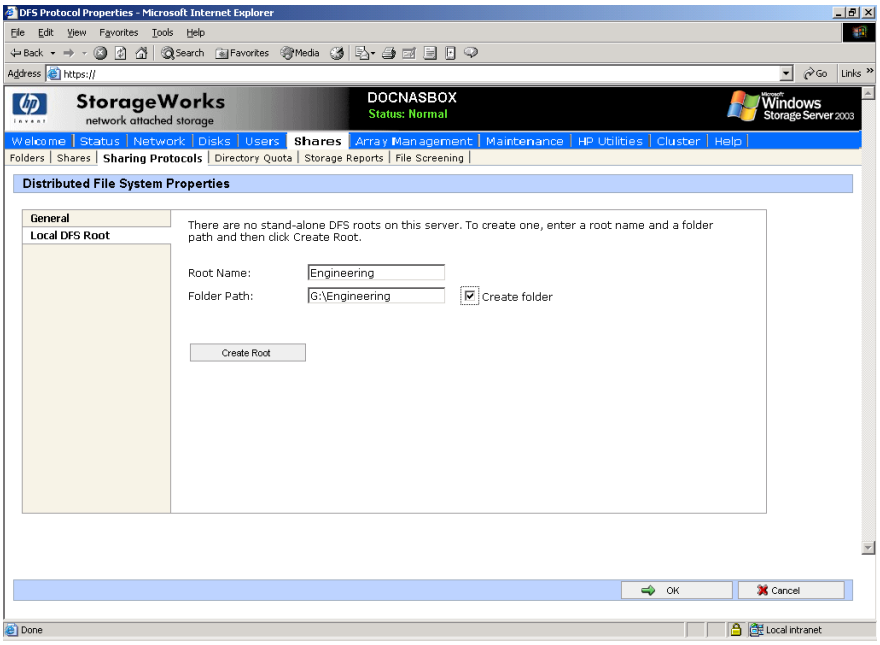

Figure 83: Local DFS Root tab

4. On the Local DFS Root tab, type the name of the DFS root in the **Root name** box.
5. In the **Folder path** box, type the path of the folder that corresponds to the root. Click **Create folder** if the folder does not exist.
6. Choose Create DFS Root, and then choose **OK**.

Deleting a Local DFS Root

The WebUI enables the deletion of a local stand-alone DFS root on the server only. The Distributed File System administrative tool must be used to manage Domain DFS Roots. Hence, if there is more than one root on the server, the first root (in alphabetical order, with local stand-alone roots grouped ahead of domain roots) will be available to be deleted. If only domain roots exist on the server, the first domain root will be listed, but it cannot be deleted using the WebUI. The WebUI can only be used to manage local stand-alone DFS roots.

To delete a local DFS root:

1. On the primary navigation bar, choose **Shares**.
2. Choose **Sharing Protocols**.
3. Select **DFS**, and then choose **Properties**. On the Local DFS Root tab, choose **Delete Root**.
4. Choose **OK**.

Publishing a New Share in DFS

Once a root has been established either on the local machine or one in the network, shares can be published in order to extend the virtual name space. For example, several shares can be created for a DFS root labeled “Engineering.” The shares might be titled “Documentation,” “Technical Specs,” and “Project Info.” When mapping to *computername**engineering*, all three of these shares would be found under the mapped drive even though they exist on different NAS devices, drives or shares points. To publish a share in a DFS root:

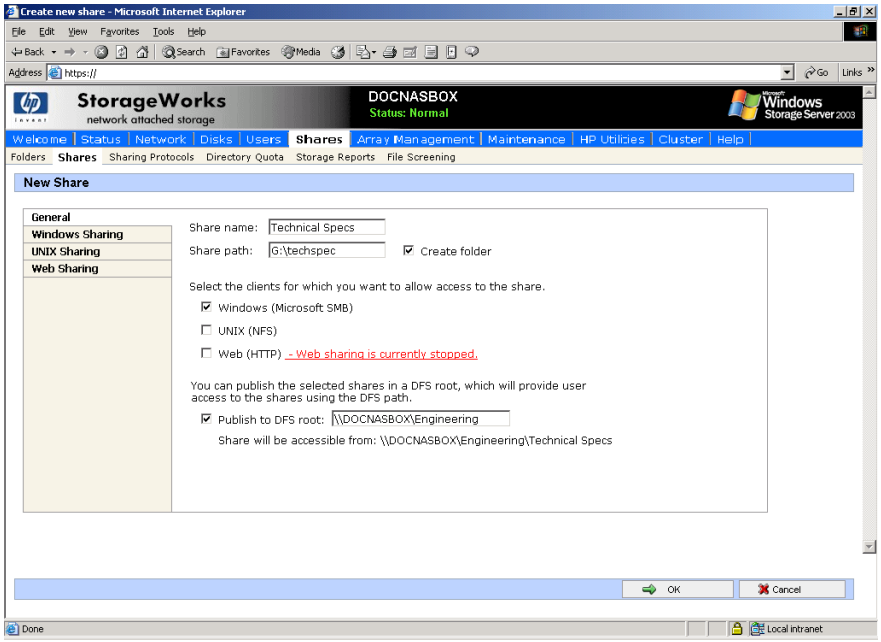

Figure 84: DFS share example

1. Select **Shares** from the WebUI.
2. Type in a new share name
3. Type in a folder name (select the checkbox **Create folder** if appropriate)
4. Verify that the Windows checkbox is selected. (DFS is dependent on the SMB protocol)
5. Under DFS, check the box if unchecked.

Note: The default behavior can be set to publish all shares to DFS. In this case the box will be checked. See the section above **Setting DFS Sharing Defaults.**

---

6. Enter in the name of the DFS root to publish the share (“Engineering” in this example). The network name will be displayed below the entry.
7. Click **OK**.

A share name will be published in the namespace.

To view the namespace, map a drive to the DFS root. All published shares will be seen in the namespace. See the example in Figure 85.

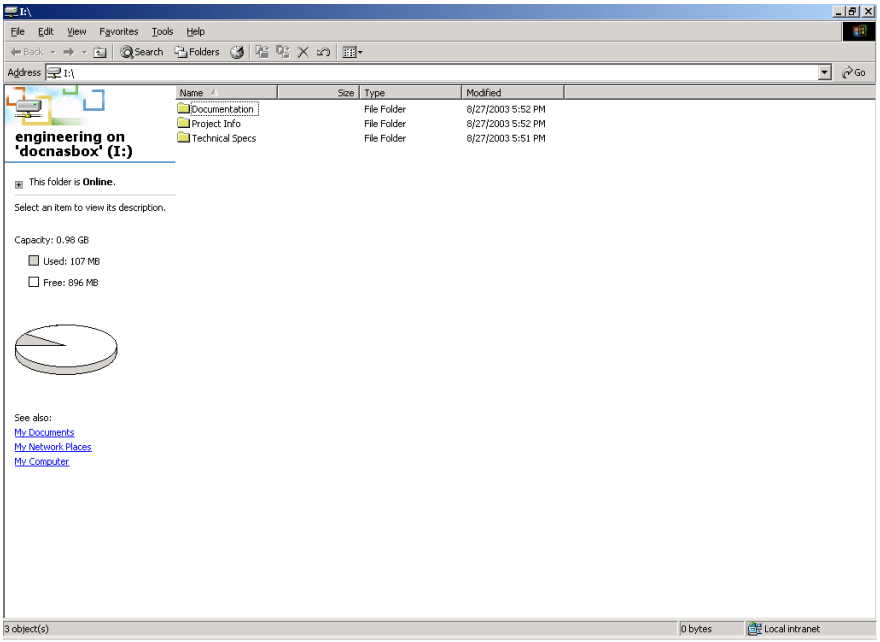

Figure 85: DFS share example, mapped drive

In this case, Documentation exists on *G:\documentation*, Technical Specs exists on *G:\technical specs* and Project Info exists on *C:\project info* on the local machine but they are all accessible via *\\DOCNASBOX\engineering*.

Publishing an Existing Share in DFS

To enable an existing shares for DFS, perform the following steps:

1. Select **Shares** from the WebUI.
2. Select the target share from the table and select **Publish in DFS**.
3. Enter the name of the DFS root to publish the share too.
4. Click **OK**.

The share will appear in the DFS underneath the DFS root.

Removing a Published Share from DFS

Once a share is published in DFS, it may be removed from the virtual namespace via the Shares Property page. To remove a share from DFS perform the following steps:

1. Select **Shares** from the WebUI.
2. Select the target share from the table and select properties.
3. Uncheck the box entitled Publish to DFS root.
4. Click **OK**.

The share will no longer appear in the DFS.

Storage Management

The storage management features built into the NAS server are composed of three main features and are applicable at the directory level of a share. These features include:

- Directory Quotas
- File Screening
- Storage Reports

Each of these feature sets are describe below. For procedures and methods, refer to the online help available within the web UI via the ? in the right hand corner of each accompanying feature management page.

Note: The storage management features are not supported in a clustered environment. In a clustered environment, these features should be uninstalled as instructed in the Cluster Installation Guide. See Figure 86.

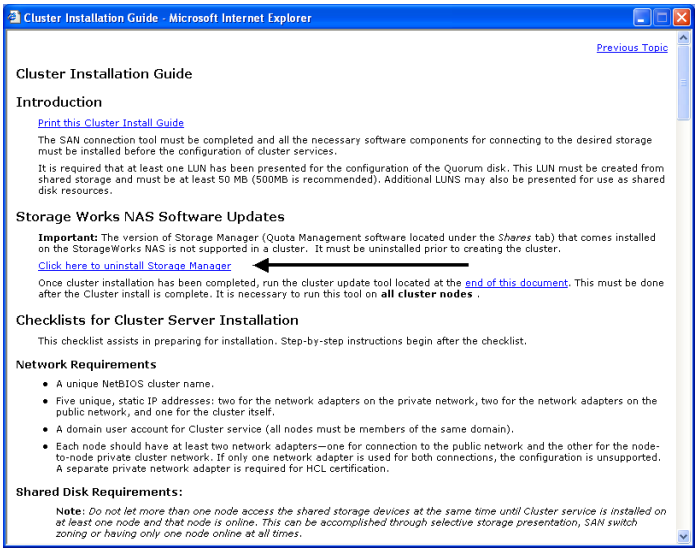

Figure 86: Uninstall storage manager

Directory Quotas

Directory quotas provide a way to limit and monitor space consumed by all files in a folder. For information on setting quotas on volumes, see Chapter 5.

Directory quotas limit the size of the managed object regardless of who writes to or who owns files in the managed object. For example, if a 50MB directory quota is set on the managed object *c:\users\JDoe*, that directory and all its contents will be limited to 50MB regardless of who owns the files in that directory or who writes files to that directory.

Directory quotas allow for the addition, deletion, monitoring, and changing of space limits for selected directories on the NAS server. Directory quotas provide disk space monitoring and control in real time, and support active and passive limits with two real-time space alarms.

The Directory Quota feature includes the following components:

- Active and passive space limits on directories
- Best practice storage resource management policies
- A severe alarm threshold
- A warning alarm threshold
- Auto discovery of drives
- Customized messages
- Alarms sent to the event log
- Alarms sent to the user
- Storage reports that can be sent to an intranet Web site
- Custom script support

The directory quota set on the system partition always has a passive limit and uses device size (capacity). If the system does not have sufficient quota to write files, it may fail. Also, if the system partition does not have enough space to write temporary files during boot, the system may not restart. Avoid this by using caution when placing quotas on the system directories.

Directory quotas use each file's allocation size to determine how much space is used. The allocation size is slightly larger than the actual space used as displayed by Windows Explorer and other Windows programs for the data in a file. This discrepancy may cause some confusion, but the Directory Quota feature is correctly charging the user for the amount of disk space actually consumed to store a file. Large cluster sizes on file allocation table (FAT) file systems may add to the confusion because the entire cluster is always allocated, regardless of the file size. NTFS file systems store very small files in the index file and typically have more reasonable cluster sizes.

Because of the differences in the amount of storage requested for a file extension operation and the amount actually allocated by Windows Storage Server 2003 for that extension, the user may be allowed to exceed his quota by as much as one cluster. For example, assume the user has a quota of 100 KB and has used 96 KB on a file system with a cluster size of 8 KB. The user creates a file of 1 KB. Windows Storage Server 2003 requests 1024 bytes be allocated for the file. Since this is less than the remaining quota for the user, the operation is allowed to continue. However, if the cluster size is 8 KB, Windows Storage Server 2003 will actually allocate 8 KB for the file. The user has now used 104 KB, and while this is allowed, future attempts to create or extend files will fail.

Establishing Directory Quotas

Directory quotas are established in a two part fashion. First a policy is defined using the policies selection from the Directories Policy Page. After a policy is established it can be assigned to a particular directory via the WebUI “New Directory Quota Wizard”. By default there are a number of predefined policies, these policies include:

- 100 MB Limit
- 500 MB Limit
- Best Practices Report
- Default
- Monitor Directory
- Partition Alert

Each of these policies provides an example of a particular policy type. Custom policies should be created to meet the needs of the environment.

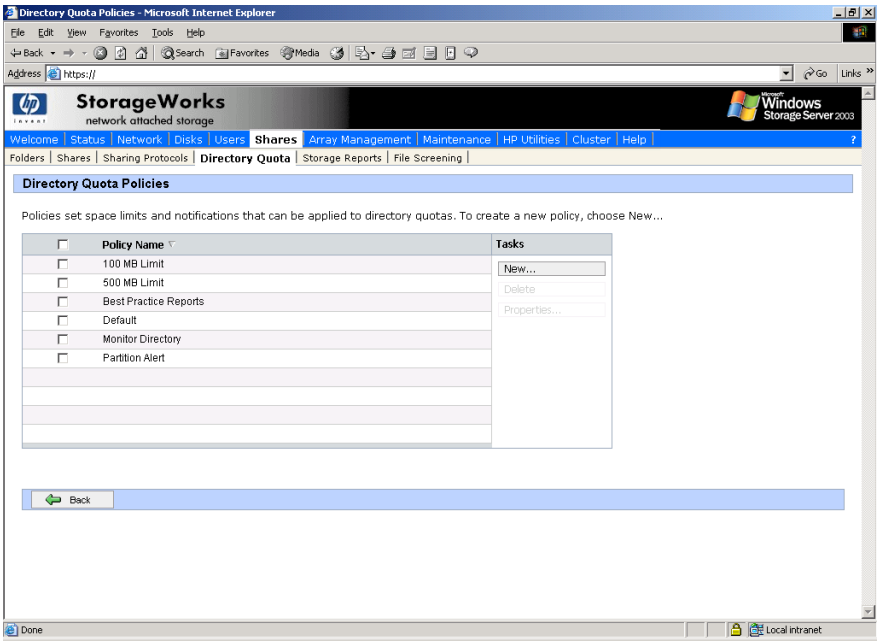

Figure 87: Policies main page

Within each policy, there are a number of configuration screens that are presented in the form of a wizard. The wizard collects the following information to create a policy:

- Name of Policy
- Disk space limit and Unit of measurement
- Passive limit (If selected the limit will issue warnings but will not prevent access.
- Alarm Threshold for severe and warning messages
- Notification for severe and warning messages

The notification field allows for the creation of a message to be sent to the eventlog of the server or via Netbios as a pop up on the client machine. It should be noted that Netbios is not supported in all customer environments and the pop up function may not be supported.The notification includes macro functionality and variable definitions for user custom messages. The help function (?) in the right hand corner of the UI provides an online guide to building these macro function messages under the topic “Directory Quota Alarm Notification.”

To modify any of these settings at a later time the properties button may be selected for a particular policy or quota. In addition, to policy settings for existing shares, default policies can be set in advance for new devices added to the system via the preferences button on the Directory Quota Page.

File Screening

File screening allows the administrator to limit or monitor files based on extension, for example disallow all .pst and .mp3 files. It should be noted that the filter is merely based on extensions and not the content of the files. Hence, if a file extension is renamed away from .mp3 for example to mpp, the filter software will allow the file to be stored. A complete online help guide in the WebUI is provided for file screening via the ? in the right hand corner of the UI.

File screening is established in the policy settings. Screening groups contain a collection of authorized and un-authorized file extensions. Filters determine which folders to exclude. Alarms, similar to the actions when a quota threshold is exceeded, can be set up when an unauthorized file type is set up.

File screening includes the following features:

- Active and passive file screening on directories
- Best practice file screening policies
- Notification alarm when file screening policy is violated
- Audit database containing screened files
- Customized alarm messages
- Alarm messages to the event log
- Alarm messages to a user
- Storage reports when alarm is activated and sent to intranet Web site
- Custom script when alarm is activated
- Real time monitoring of file screening activity

Use caution when placing screening parameters on the system partition. If certain classes of files are screened from the system partition, the operating system may not have the access to save temporary working files. It is a good idea to exclude systems directories from screening. Another option is to create a passive screening policy that allows files to be saved but the file activity to be logged.

File Screening essentially has the same feature sets as directory quotas with one exception. Groupings of file types are first created, such as multimedia files, graphics, etc. These groups are then placed in a particular policy. A file screen is then enabled on a directory and the various policies are applied for a particular directory. Lastly, the same types of alert notification is allowed as in the case of the directory quotas. See the online help for additional information.

Storage Reports

Storage reports allow the administrator to analyze the contents of the storage server via standard reports for common tasks. The reports can be displayed using text, simple HTML tables, or Active HTML. When using Active HTML, the ActiveX control provides graphs. A complete online help guide in the WebUI is provided for reporting via the ? in the right hand corner of the UI.

Reports can be scheduled, or produced on demand.

Storage reports address disk usage, wasted space, file ownership, security, and administration. Reports can run interactively, schedule on a regular basis, or run as part of a storage resource management policy when disk space utilization reaches a critical level.

Storage reports may be presented in Hyper Text Markup Language (HTML) and text (.txt) formats. The output formats can be e-mailed to a list of users.

The following features are included with storage reports:

- Best practice storage resource management reports
- Integration with best practice storage resource management policies
- Scheduled storage reports
- Storage reports sent to an intranet Web site
- Storage reports sent to users through e-mail

Note: Large storage reports should be scheduled for off-hours.

Print Services

Printer services are a new feature added to the NAS server that has not been available previously. The new service supports network printers only and is not intended for use with locally attached printers (USB or Parallel port connected).

Note: See the Cluster Administration chapter for information on clustering a print spooler.

If the NAS server is a part of an Active Directory domain vs Workgroup, the NAS server enables the following management features:

- Restrict access to a printer based domain user accounts
- Publish shared printers to Active Directory to aid in search for the resource

Before adding a print server role the following check list of items should be followed:

1. **Determine the operating system version of the clients that will send jobs to this printer**. This information is used to select the correct client printer drivers for the client and server computers utilizing the printer. Enabling this role on the print server allows for the automatic distribution of these drivers to the clients. Additionally, the set of client operating systems determines which of these drivers need to be installed on the server during the print server role installation.
2. **At the printer, print a configuration or test page that includes manufacturer, model, language, and installed options**. This information is needed to choose the correct printer driver. The manufacturer and model are usually enough to uniquely identify the printer and its language. However, some printers support multiple languages, and the configuration printout usually lists them. Also, the configuration printout often lists installed options, such as extra memory, paper trays, envelope feeders, and duplex units.
3. **Choose a printer name**. Users running Windows-based client computers choose a printer by using the printer name. The wizard that you will use to configure your print server provides a default name, consisting of the printer manufacturer and model. The printer name is usually less than 31 characters in length.
4. **Choose a share name**. A user can connect to a shared printer by typing this name, or by selecting it from a list of share names. The share name is usually less than 8 characters in length for compatibility with MS-DOS and Windows 3.x clients.
5. (Optional) **Choose a location description and a comment**. These can help identify the location of the printer and provide additional information. For example, the location could be “Second floor, copy room” and the comment could be “Additional toner cartridges are available in the supply room on floor 1."

Configuring the Print Server

To set up a print server:

1. Click **Start > Settings > Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.
2. Click on **Add or Remove a Roll**.
3. A wizard will start. Click **Next**.
4. Select Printer Server from the list of Server Roles and click **Next**.
5. Select Windows 2000 and Windows XP clients only and click **Next**.

Note: While the "All Windows" support may be selected at this step, it is more efficient to add the alternative operating systems on each printer after the wizards are complete. See section below on "Adding Additional Operating System Support".

6. Click **Next on the Summary page** and an Add Printer Wizard will start.
7. Select Local Printer and uncheck "automatically detect install my plug and play printers" click **Next**.

Note: Local Printer is used to create a TCP/IP port connections to a network enabled printer over the network. The NAS server only supports network attached printers and does not support directly connected printers via USB or Parallel Port.

8. Select **Create a new port**, and select **Standard TCP/IP Port** (recommended).
9. The Add Standard TCP/IP Printer Port Wizard starts. Click **Next**.
10. Type the name or IP address of the printer. The IP address is usually listed on the printer configuration page. The wizard completes the Port Name field. Click **Next**.
11. The wizard attempts to connect to the printer. If the wizard is able to connect, the **Completing the Add Standard TCP/IP Printer Port** Wizard page appears, and click **Finish**. If the wizard is not able to connect, the **Additional Port Information Required** page appears.
    a. Verify that the ip address or name that was entered is correct.
    b. Select **Standard** to identify the printer network adapter. A list of manufacturers and models of the network adapters will be displayed. Select the appropriate printer from the Standard list.
    c. If the printer network adapter uses nonstandard settings, click **Custom** and then click **Settings**. The **Configure Standard TCP/IP Port Monitor** page appears. Specify the settings that are recommended by the manufacturer of the printer network adapter, and then click **OK**.
    d. Click **Next**.
12. Select the manufacturer and the type of printer from the presented list, click **Next**. If the printer does not exist in the list, click have disk and load the drivers or select a compatible driver.
13. Enter the name of the desired printer to be presented on the NAS device, click **Next**.
14. Enter a Share Name for the printer that will used on the network, click **Next**.
15. Enter a location description and a comment, click **Next**.
16. Select Print a test page and click **Next**.
17. Uncheck the restart the add printer wizard if adding only one printer, click **Finish**.
18. A test page will printer, click ok if the page printed otherwise select troubleshoot.
19. If the restart the add printer wizard was selected the wizard will restart to add an additional printer. Repeat the steps above for adding an additional printer.

Removing the Print Server Role

To remove the print server role:

1. Click **Start > Settings > Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.
2. Click on **Add or Remove a Roll**.
3. A wizard will start. Click **Next**.
4. Select **Printer Server** from the list of Server Roles and click **Next**.
5. Select the checkbox **Remove the printer role**, click **Next**.
6. The Printer role will be removed, click **Finish**.

Adding an Additional Printer

To add additional printers to the NAS device:

1. Select **Start > Settings > Printers and Faxes > Add Printer**.
2. The add printer wizard will start. Click **Next**.
3. Select Local Printer and uncheck “automatically detect install my plug and play printers.” Click **Next**.

Note: Local Printer is used to create a TCP/IP port connections to a network enabled printer over the network. The NAS server only supports network attached printers and does not support directly connected printers via USB or Parallel Port.

---

4. Select **Create a new port**, and select **Standard TCP/IP Port** (recommended).
5. The **Add Standard TCP/IP Printer Port** Wizard starts. Click **Next**.
6. Type the name or IP address of the printer. The IP address is usually listed on the printer configuration page. The wizard completes the Port Name field. Click **Next**.
7. The wizard attempts to connect to the printer. If the wizard is able to connect, the Completing the **Add Standard TCP/IP Printer Port** Wizard page appears, and click **Finish**. If the wizard is not able to connect, the **Additional Port Information Required** page appears.
   a. Verify that the IP address or name that was entered is correct.
   b. Select **Standard** to identify the printer network adapter. A list of manufacturers and models of the network adapters is displayed. Select the appropriate printer from the Standard list.
   c. If the printer network adapter uses nonstandard settings, click **Custom** and then click **Settings**. The **Configure Standard TCP/IP Port Monitor** page appears. Specify the settings that are recommended by the manufacturer of the printer network adapter, and then click **OK**.
   d. Click **Next**.
8. Select the manufacturer and the type of printer from the presented list, click **Next**. If the printer does not exist in the list, click **have disk** and load the drivers or select a compatible driver.
9. Enter the name of the desired printer to be presented on the NAS device, click **Next**.

10. Enter a Share Name for the printer that will used on the network, click **Next**.
11. Enter a location description and a comment, click **Next**.
12. Select **Print a test page** and click **Next**.
13. Click **Finish**. A test page prints. Click **OK** if the page printed otherwise select **Troubleshoot**.

Adding Additional Operating System Support

By default, support is added for Windows 2000 and Windows XP. If the client base is composed of other Windows operating systems, additional printer drivers will need to be loaded. To load an additional driver for client download:

1. Select **Start > Settings > Printers and Faxes**, right-click on the printer to manage.
2. Select **Properties**.
3. Select the **Sharing** tab.
4. Select **Additional Drivers**.
5. Select the desired operating systems and click **OK**.
6. A dialog will appear to add the additional drivers from disk.

Installing Print Services for UNIX

1. Log on as administrator or as a member of the Administrators group.
2. Click **Start > Control Panel**, and then click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components**.
4. In the **Components** list, click **Other Network File and Print Services** (but do not select or clear the check box), and then click **Details**.
5. In the Subcomponents of **Other Network File and Print Services** list, click to select **Print Services for UNIX**, if appropriate to the print services that you want to install:

   Print Services for UNIX: This option permits UNIX clients to print to any printer that is available to the print server.

Note: When you install Print Services for UNIX, this automatically installs the LPR port and the TCP/IP Print Server service.

---

6. Click **OK**, and then click **Next**.
7. Click **Finish**.

Microsoft Services for NFS

8

Microsoft Services for NFS is a comprehensive software package designed to provide complete UNIX environment integration into a Windows NT, Windows 2000, Windows Storage Server 2003, or Active Directory domain file server. Services for NFS manages tasks on both Windows and UNIX platforms. Tasks include creating NFS exports from Windows and administering user name mappings.

The following Services for NFS components are included in the NAS server:

- Server for NFS
- User Name Mapping
- NFS Authentication

Note: Services for NFS can be implemented in both clustered and non-clustered environments. This chapter discusses Services for NFS in a non-clustered deployment. For additional information that is specific to a cluster, see the Cluster Administration chapter.

Server for NFS

Services for NFS enables UNIX clients to access a file share on the NAS server. The Services for NFS server supports NFS Version 2 and Version 3, and supports them both on the TCP and UDP network protocols.

Services for NFS is more fully integrated into the operating system than other third-party NFS server packages. The administrative interface for NFS exports is similar to the Server Message Block (SMB) sharing interface used by Windows platforms. With Server for NFS properly configured, the administrator can create shares that are simultaneously accessible by multiple client types. For example, some of the options for shares include configurations for CIFS/SMB sharing only, simultaneous NFS/CIFS/SMB sharing, simultaneous NFS/CIFS/SMB/HTTP sharing, or simply NFS only sharing.

Authenticating User Access

NFS export access is granted or denied to clients based on client name or IP address. The server determines whether a specific client machine has access to an NFS export. No user logon to the NFS server takes place when a file system is exported by the NFS server. Permission to read or write to the export is granted to specific client machines. For example, if client machine M1 is granted access to an export but client M2 is not, user jdoe can access the export from M1 but not from M2.

Permissions are granted on a per-export basis; each export has its own permissions, independent of other exports on the system. For example, file system a can be exported to allow only the Accounting department access, and file system m can be exported allowing only the Management department access. If a user in Management needs access to the Accounting information, the A export permissions can be modified to let that one user's client machine have access. This modification does not affect other client access to the same export, nor does it allow the Management user or client access to other exports.

After the client machine has permission to the export, the user logon affects file access. The client machine presents the UNIX user ID (UID) and group ID (GID) to the server. When the computer accesses a file, the UID and GID of the client are transferred to a Windows user ID and group ID by the mapping server. The ACLs of the file or directory object being requested are then compared against the mapped Windows login or group ID to determine whether the access attempt should be granted.

Note: User credentials are not questioned or verified by the NFS server. The server accepts the presented credentials as valid and correct.

If the NFS server does not have a corresponding UID or GID, or if the administrator has set other conditions to filter out the user, a process called squashing takes effect. Squashing is the conversion of an unknown or filtered user to an anonymous user. This anonymous user has very restricted permissions on the system. Squashing helps administrators manage access to their exports by allowing them to restrict access to certain individuals or groups and to squash all others down to restricted (or no) access. Squashing enables the administrator to allow permissions instead of denying access to all the individuals who are not supposed to have access. See “NFS User and Group Mappings” later in this chapter for specific information about creating and maintaining mappings.

Indicating the Computer to Use for the NFS User Mapping Server

During the processes of starting and installing the NAS server, the name localhost is assigned by default to the computer. It is assumed that the NAS server is the computer that will be used for user name mapping.

If there are other mapping servers and a machine other than the localhost that will store user name mappings, the name of that computer must be indicated, as detailed below:

1. Use **Remote Desktop** to access the **NAS Management Console**, click **File Sharing**, **Microsoft Services for Network File System**. Click **Settings**. Figure 88 is an example of the Server for NFS user interface.
2. In the **Computer** name box of the user-mapping screen, type the name of the computer designated for user mapping and authentication.
3. Localhost is the computer name assigned by default on the NAS server. To control user mapping from a different computer, enter the name of that computer.

Note: If a machine other than the localhost is to be used, make sure that the user name mapping service is installed and running on that machine.

Note: If the authentication software is not installed on all domain controllers that have user name mappings, including Primary Domain Controllers, Backup Domain Controllers, and Active Directory Domains, then domain user name mappings will not work correctly.

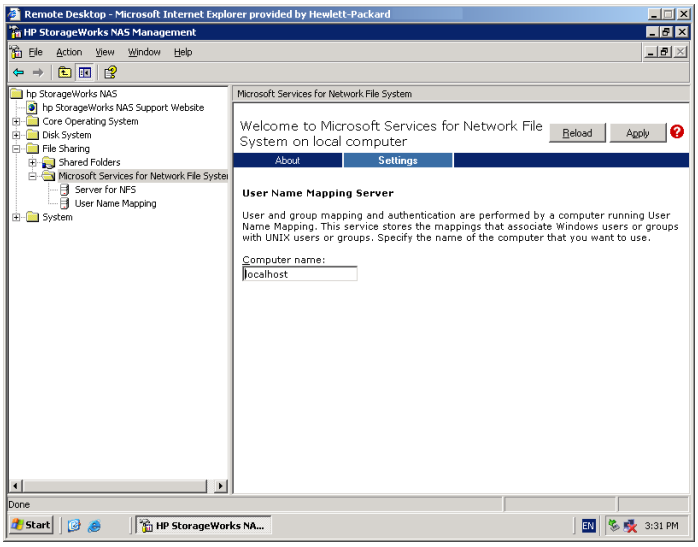

Figure 88: Microsoft Services for NFS screen, Settings tab

Logging Events

Various levels of auditing are available. Auditing sends Services for NFS events to a file for later review and establishes log-setting behavior. Some behavior examples include events logged and log file size. See the online Services for NFS help for more information.

1. Use Remote Desktop to access the NAS Management Console, click **File Sharing**, **Microsoft Services for Network File System**, **Server for NFS**. Click the **Logging** tab.
2. To log events to the event viewer application log, click the check box for **Log events to event log**.
3. To log selected event types, click the check box for **Log events in this file** on the screen.
4. Enter a filename or use the default filename provided (*rootdrive\MSNFS\log\nfssvr.log*) and log file size (7 MB default). The default log file is created when the changes are applied.

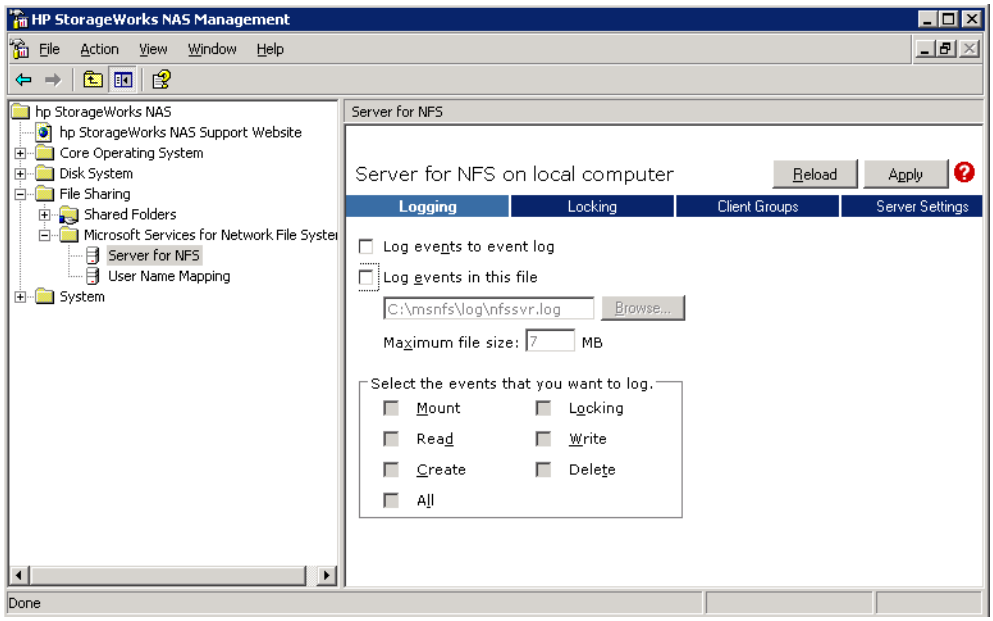

Figure 89: Server for NFS screen, Logging tab

Server for NFS Server Settings

The NAS server has new features for Services for NFS included in the Services for NFS administration GUI. The new features include settings that affect performance, such as toggling between TCP and UDP NFS versions 2 and 3. Other Server for NFS server settings include those that affect how file names are presented to NFS clients, such as allowing hidden files and allowing case sensitive lookups.

Note: The NFS Server service needs to be restarted when changing these settings. Notify users when stopping and restarting the NFS Server service.

Use Remote Desktop to access the NAS Management Console. Choose **File Sharing**, **Microsoft Services for Network File System**. Click **Server for NFS**, then **Server Settings**.

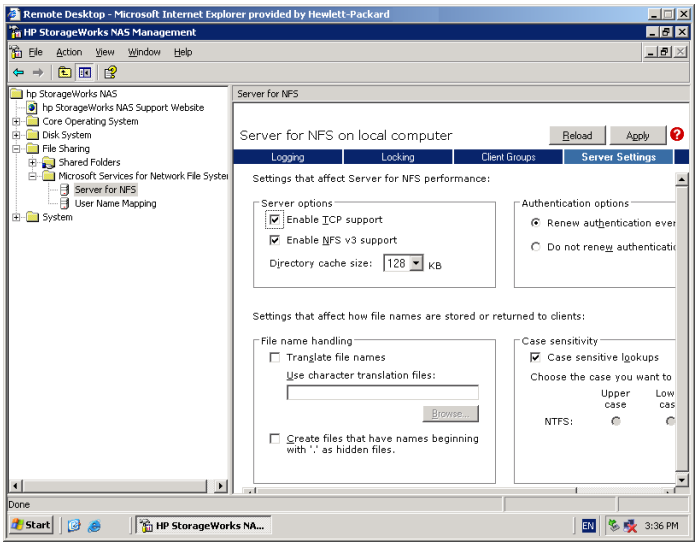

Figure 90: Server for NFS screen, Server Settings tab

Installing NFS Authentication Software on the Domain Controllers and Active Directory Domain Controllers

The NFS Authentication software must be installed on all Primary Domain Controllers (PDCs) and backup domain controllers (BDCs) that have Windows users mapped to UNIX users. This includes Active Directory domains. For instructions on setting up user mappings, see "NFS User and Group Mappings."

Note: If the authentication software is not installed on all domain controllers that have user name mappings, including Primary Domain Controllers, Backup Domain Controllers, and Active Directory Domains, then domain user name mappings will not work correctly.

The Authentication software can be installed in two different ways— the Share method or CD method.

To install the Authentication software on the domain controllers (Share Method):

1. Share out *C:\hpnas\components\SFU 3.0*.
2. Locate the *setup.exe* file located in the *SFU 3.0* directory of the NAS server.
3. On the domain controller where the Authentication software is being installed use Windows Explorer to:
   a. Open the shared directory containing *setup.exe*.
   b. Double-click the file to open it. Windows Installer is opened.

Note: If the domain controller used does not have Windows Installer installed, locate the file InstMSI.exe on the SFU 3.0 directory and run it. After this installation, the Windows Installer program starts when opening setup.exe.

---

4. In the Microsoft Windows Services for UNIX Setup Wizard dialog box, click **Next**.
5. In the User name box, type your name. If the name of your organization does not appear in the Organization box, type the name of your organization there.
6. In the CD Key boxes, type the product key found on the back of the CD-ROM case that is included with the NAS server software, and then click **Next**.
7. Read the End User License Agreement carefully. If you accept the terms of the agreement, click I accept the terms in the License Agreement, and then click Next to continue installation. If you click I do not accept the License Agreement (Exit Setup), the installation procedure terminates.
8. Click Custom Installation, and then click **Next**.
9. In the Components pane, click the down arrow next to Windows Services for UNIX, and then click **Entire component will not be available**.
10. Click the plus sign (+) next to Authentication Tools.
11. In the Components pane, click the plus sign (+) next to Authentication Tools.
12. Click **Server for NFS Authentication**, click **Will be installed on local hard drive**, and then click **Next**.
13. Follow the remaining instructions in the Wizard.

Microsoft Services for Unix 3.0 CD has been included with the NAS server and is needed for the following procedure.

To install the Authentication software on the domain controllers (CD Method):

1. Insert the Microsoft Windows Services for UNIX compact disc into the CD-ROM drive of the domain controller.
2. In the Microsoft Windows Services for UNIX Setup Wizard dialog box, click **Next**.
3. In the User name box, type your name. If the name of your organization does not appear in the Organization box, type the name of your organization there.
4. In the CD Key boxes, type the product key found on the back of the CD-ROM case, and then click **Next**.
5. Read the End User License Agreement carefully. If you accept the terms of the agreement, click **I accept the terms in the License Agreement**, and then click **Next** to continue installation. If you click **I do not accept the License Agreement** (Exit Setup), the installation procedure terminates.
6. Click **Custom Installation**, and then click **Next**.
7. In the Components pane, click the down arrow next to Windows Services for UNIX, and then click **Entire component will not be available**.
8. Click the plus sign (+) next to Authentication Tools.
9. In the Components pane, click the plus sign (+) next to Authentication Tools.
10. Click **Server for NFS Authentication**, click **Will be installed on local hard drive**, and then click **Next**.
11. Follow the remaining instructions in the Wizard.

Note: NFS users can be authenticated using either Windows domain accounts or local accounts on the Windows server. Server for NFS Authentication must be installed on all domain controllers in the domain if NFS users will be authenticated using domain accounts. Server for NFS Authentication is always installed on the 9000s (or 4000s depending), which also runs the Server for NFS.

Understanding NTFS and UNIX Permissions

When creating a NFS export, make sure that the NTFS permissions on the share allows the correct permissions that you want assigned for users/groups. The following will help clarify the translation between Unix and NTFS permissions:

- The UNIX read bit is represented within NTFS as the List Folder/Read Data permission
- The UNIX write bit is represented within NTFS as the Create File/Write Data, Create Folders/Append Data, Write Attributes, and Delete Subfolders and Files permissions
- The UNIX execute bit is represented within NTFS as the Traverse Folder/Execute File permission

NFS File Shares

NFS file shares are created in the same manner as other file shares, however there are some unique settings. Procedures for creating and managing NFS file shares are documented in the same sections as creating file shares for other protocols. See the "Folder and Share Management" chapter for more information.

Note: NFS specific information is extracted from the "Folder and Share Management" chapter and duplicated below.

Complete share management is performed through the **Shares** menu option of the WebUI. Tasks include:

- Creating a new share
- Deleting a share
- Modifying share properties

Each of these tasks is discussed in this section.

Creating a New Share

To create a new NFS file share:

1. From the WebUI main menu, select the **Shares** tab and then select the **Shares** option. The **Shares** dialog box is displayed. From the **Shares** dialog box, click **New**. The **General** tab of the **Create a New Share** dialog box is displayed.

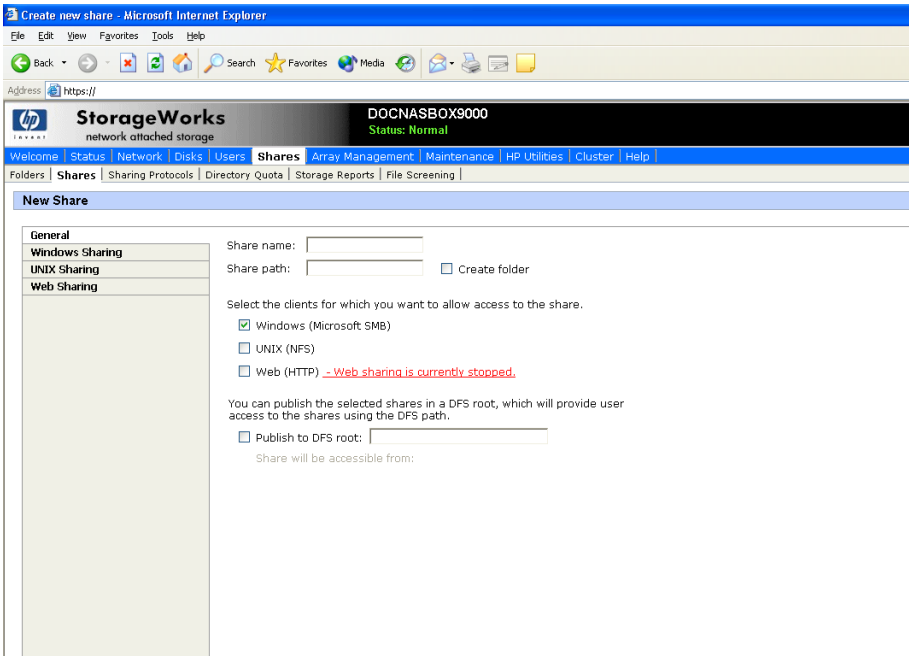

Figure 91: Create a New Share dialog box, General tab

2. In the **General** tab, enter the share name and path. Check the **Unix (NFS)** client protocol check box.

Note: Uncheck the Microsoft SMB option if you do not want to allow SMB access to the share.

Note: NFS service does not support the use of spaces in the names for NFS file shares. NFS translates any spaces in an export into an underscore character. If you plan to use the same name when sharing a folder through SMB, and then exporting it through NFS, do not put spaces in the SMB share name.

   To create a folder for the share, check the indicated box and the system will create the folder at the same time it creates the share.

3. Select the **NFS Sharing** tab to enter NFS specific information. See "Modifying Share Properties" for information on this tab.
4. After all share information is entered, click **OK**.

The default NFS share properties are **All Machines read only with root and anonymous access disallowed**. See the section, "Modifying Share Properties" in this chapter to change the default permissions.

Deleting a Share

Caution: Before deleting a share, warn all users to exit that share. Then confirm that no one is using the share.

To delete a share:

1. From the **Shares** menu, select the share to be deleted, and then click **Delete**.
2. Verify that this is the correct share, and then click **OK**.

Modifying Share Properties

To change share settings:

1. From the **Shares** menu, select the share to modify and then click **Properties**. The **General** tab of the **Share Properties** dialog box is displayed.

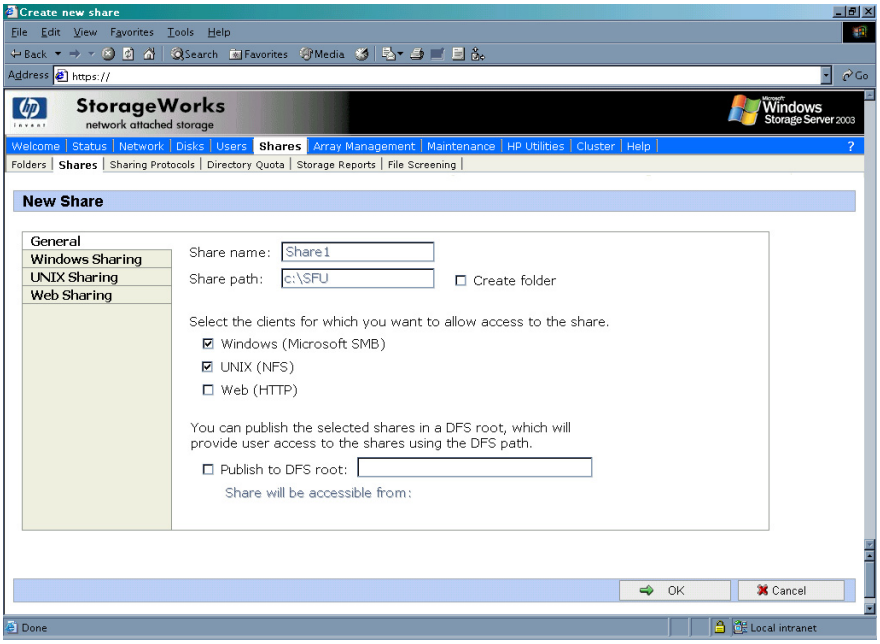

Figure 92: Share Properties dialog box, General tab

The name and path of the selected share is displayed.

2. To enter or change client protocol information, check the **UNIX (NFS)** client type box and then click the **UNIX Sharing** tab.

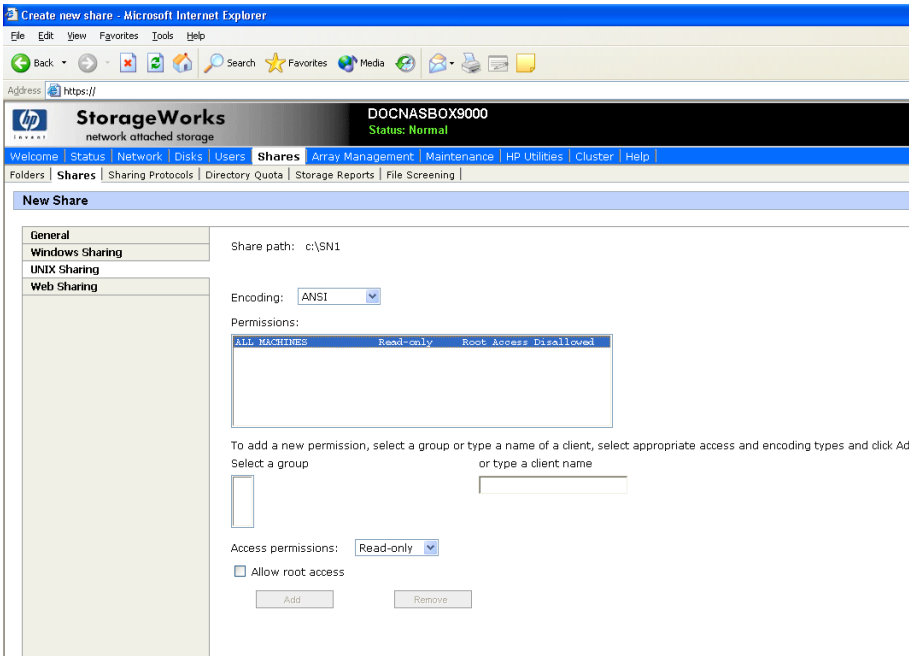

Figure 93: UNIX Sharing tab

3. From the **UNIX Sharing** tab of the **Share Properties** dialog box,
   a. Indicate the allowed clients.

      Select the machine to include in the **Select a group** box or manually enter the NFS client computer name or IP address. Then click **Add**.

   b. Indicate the access permissions.

      Select the machine from the main user display box and then select the appropriate access method from the **Access permissions** drop down box.

      The types of access are:

      - **Read-only**—Use this permission to restrict write access to the share.
      - **Read-write**—Use this permission to allow clients to read or write to the share.
      - **No access**—Use this permission to restrict all access to the share.

4. Select whether or not to allow root access. Check the **Allow root access** checkbox to add the root permission.

   - **Read-only + Root**—Use this permission to restrict write access to the share. Use this permission to assign administrative access to the share. This will allow the client computer to have root access to the NFS share. Map the UNIX root user to the Windows user Administrator. Also, map the group that this UNIX root belongs to, to the Windows group Administrator.
   - **Read-write + Root**—Use this permission to allow clients to read or write to the share. Use this permission to assign administrative access to the share. This will allow the client computer to have root access to the NFS share. Map the UNIX root user to the Windows user Administrator. Also, map the group that this UNIX root belongs to, to the Windows group Administrator.

5. After all UNIX sharing information is entered, click **OK**.

Anonymous Access to an NFS Share

It may be desirable to add anonymous access to a share. An instance would be when it is not desirable or possible to create and map a UNIX account for every Windows user. A UNIX user whose account is not mapped to a Windows account is treated by Server for NFS as an anonymous user. By default, the user indentifier (UID) and group identifier (GID) is -2.

For example, if files are created on an NFS Share by UNIX users whose are not mapped to Windows users, the owner of those files will be listed as anonymous user and anonymous group, (-2,-2).

By default, Server for NFS does not allow anonymous users to access a shared directory. When an NFS share is created, the anonymous access option can be added to the NFS share. The values can be changed from the default anonymous UID and GID values to the UID and GID of any valid UNIX user and group accounts.

When allowing anonymous access to an NFS Share, the following must be performed by a user with administrative privileges due to Windows Storage Server 2003 security with anonymous users and the Everyone group.

1. From the WebUI, select **Maintenance**.
2. Click **Remote Desktop**. Log on to the NAS machine.
3. Click **Start** > **Control Panel** > **Administrative Tools**, and then click Local Security Policy.
4. In Security Settings, double-click Local Policies, and then click Security Options.
5. Right-click “Network access: Let Everyone permissions apply to anonymous users,” and then click Properties.
6. To allow permissions applied to the Everyone group to apply to anonymous users, click Enabled. The default is Disabled.
7. The NFS server service will need to be restarted. From a command prompt, type `net stop nfssvc`. Then type `net start nfssvc`. Notify users before restarting the NFS service.
8. Assign the Everyone group the appropriate permissions on the NFS Share.
9. Enable anonymous access to the share.

To enable anonymous access to an NFS share, do the following.

1. Open Windows Explorer by clicking **Start** > **Run**, and typing `explorer`.
2. Navigate to the NFS share.
3. Right-click the NFS Share and click **Properties**.
4. Click **NFS Sharing**.
5. Click the checkbox next to Allow Anonymous Access.
6. Change from the default of -2,-2 if desired.
7. Click **Apply**.
8. Click **OK**.

Encoding Types

Encoding types can be selected using the WebUI. These include the default ANSI as well as EUC-JP. Other encoding types include:

- ANSI (default)
- BIG5 (Chinese)
- EUC-JP (Japanese)
- EUC-KR (Korean)
- EUC-TW (Chinese)
- GB2312-80 (Simplified Chinese)
- KSC5601 (Korean)
- SHIFT-JIS (Japanese)

If the option is set to ANSI on systems configured for non-English locales, the encoding scheme is set to the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales:

- Japanese: SHIFT-JIS
- Korean: KS C 5601-1987
- Simplified Chinese: GB
- Traditional Chinese: BIG5

NFS Only

Microsoft Services for NFS allows the option of setting up NFS Shares for NFS access only.

The NFS Only option provides faster NFS performance and is intended for NFS clients only. The executable file, *nfsonly.exe*, allows a share to be modified to do more aggressive caching to improve NFS performance. This option can be set on a share-by-share basis. Do not use this function on any file share that can be accessed by any means other than by NFS clients, as data corruption can occur.

The syntax of this command is:

```
nfsonly <sharename> [/enable|disable]
```

- Sharename is the name of the NFS share
- The /enable option turns on NfsOnly for the specified share
- The /disable option turns off NfsOnly for the specified share

The NFS service must be restarted after setting up an NFS Only share. Notify users when the NFS service is restarted.

NFS Protocol Properties Settings

Parameter settings for the NFS protocol are entered and maintained through the WebUI in the **NFS Properties** dialog box. To access the **NFS Properties** dialog box, select **Shares**, **Sharing Protocols**. Then, select the **NFS Protocol** radio button and click **Properties**.

The **NFS Properties** menu is displayed.

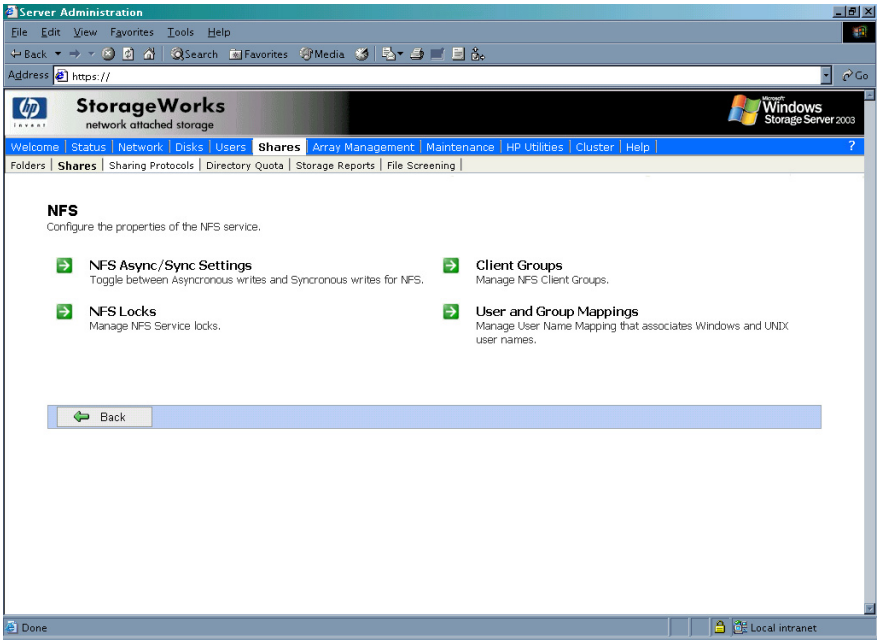

Figure 94: NFS Sharing Protocols menu

NFS properties include:

- Async/Sync Settings
- NFS Locks
- Client Groups
- User and Group Mappings

Settings for asynchronous/synchronous writes and service locks are discussed together in the following paragraphs of this chapter.

Client groups and user and group mappings are each discussed in separate sections later in this chapter.

NFS Async/Sync Settings

As mentioned in a previous section, there are two versions of NFS: Version 2 and Version 3. Version 3 supports additional file operations that Version 2 did not have, such as asynchronous file operations.

To indicate whether to use asynchronous or synchronous write settings:

1. From the WebUI, access the **NFS Protocol Properties** menu by selecting **Shares**, **Sharing Protocols**. Select **NFS Protocol** and then click **Properties**. The **NFS Properties** menu is displayed.
2. In the **NFS Properties** menu, select **NFS Async/Sync Settings**. The **NFS Async/Sync Settings** dialog box is displayed.
3. Select the desired write setting. The default setting is Synchronous writes.

Note: Using synchronous writes allows for greater data integrity. Asynchronous writes will increase performance but will reduce data integrity as the data is cached before being written to disk. Changing the write state causes the NFS service to be restarted. Notify users before toggling this setting.

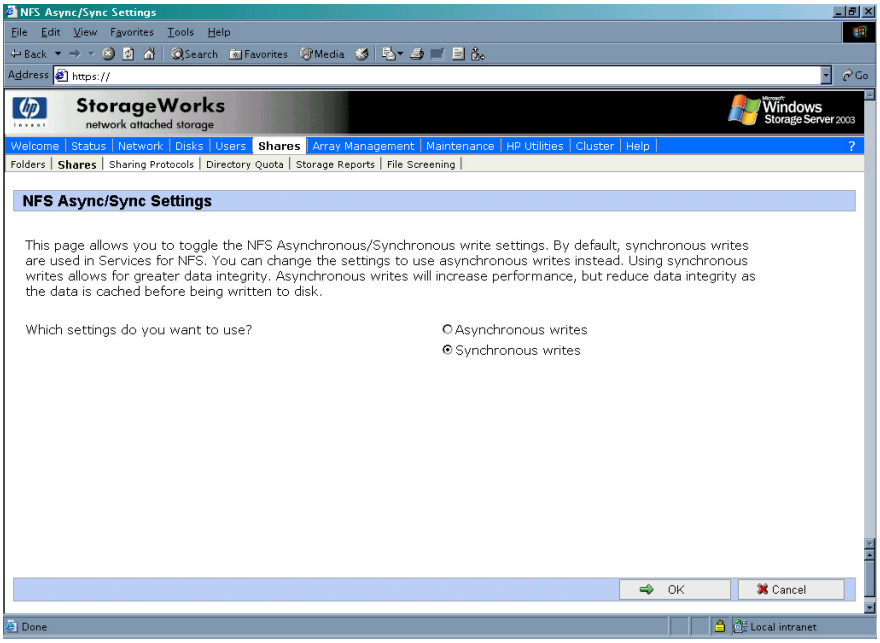

Figure 95: NFS Async/Sync Settings dialog box

NFS Locks

NFS supports the ability to lock files. File locking helps prevent two or more users from working with the same files at the same time.

NFS locking depends on the software application components to manage the locks. If an application does not lock a file or if a second application does not check for locks before writing to the file, nothing prevents the users from overwriting files.

To enter locking parameters:

1. From the WebUI, access the **NFS Protocol Properties** menu by selecting **Shares**, **Sharing Protocols**. Select **NFS Protocol** and then click **Properties**.

   The **NFS Properties** menu is displayed.

2. In the **NFS Properties** menu, select **NFS Locks**. The **NFS Locks** dialog box is displayed. Figure 96 is an illustration of the **NFS Locks** dialog box.

   All clients that have locks on system files are listed in the **Clients that hold locks** box.

3. To manually clear locks that a client has on files, select the client from the displayed list, and then click **OK**.

4. To indicate the amount of time after a system failure that the locks are kept active, enter the number of seconds in the **Wait period** box.

The NAS server keeps the locks active for the specified number of seconds, while querying the client to see if it wants to keep the lock. If the client responds within this time frame, the lock is kept active. Otherwise, the lock is cleared.

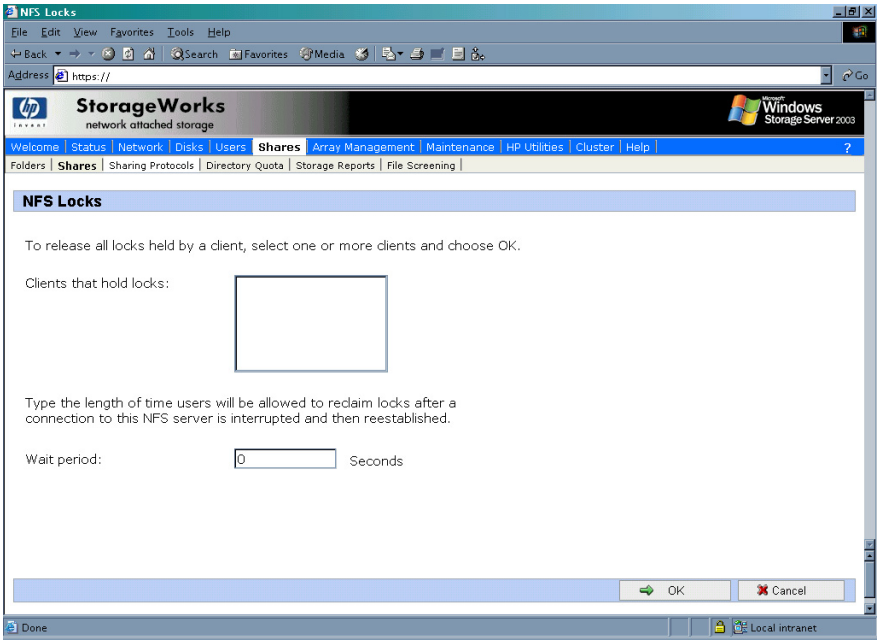

Figure 96: NFS Locks dialog box

NFS Client Groups

The Client Groups feature gives administrators a method of assigning access permissions to a set of clients. The administrator creates a client group, gives it a name, and then inserts clients into the group by client name or IP address. After the client group is created, the administrator adds or removes permissions for the entire group, instead of allowing or denying access for each individual client machine.

Proper planning includes control over the naming conventions of client groups and users. If the client group is given the same name as a client, the client is obscured from the view of the server. For example, assume that a client d4 exists. If a client group called d4 is created, permissions can no longer be assigned to just the client d4. Any reference to d4 now refers to client group d4.

To manage NFS client groups:

1. From the WebUI, access the **NFS Protocol Properties** dialog box by selecting **Shares**, **Sharing Protocols**. Select **Client Groups**. The **NFS Client Groups** dialog box is displayed.

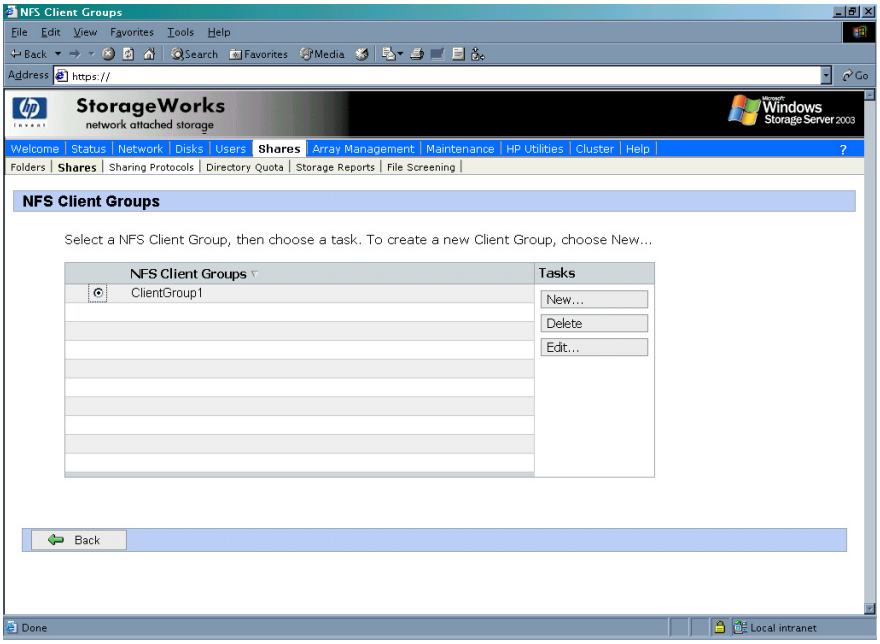

Figure 97: NFS Client Groups dialog box

The following tasks are available:

- Adding a new client group
- Deleting a client group
- Editing client group information

Adding a New Client Group

To add a new client group:

1. From the **NFS Client Groups** dialog box, click **New**. The **New NFS Client Group** dialog box is displayed.

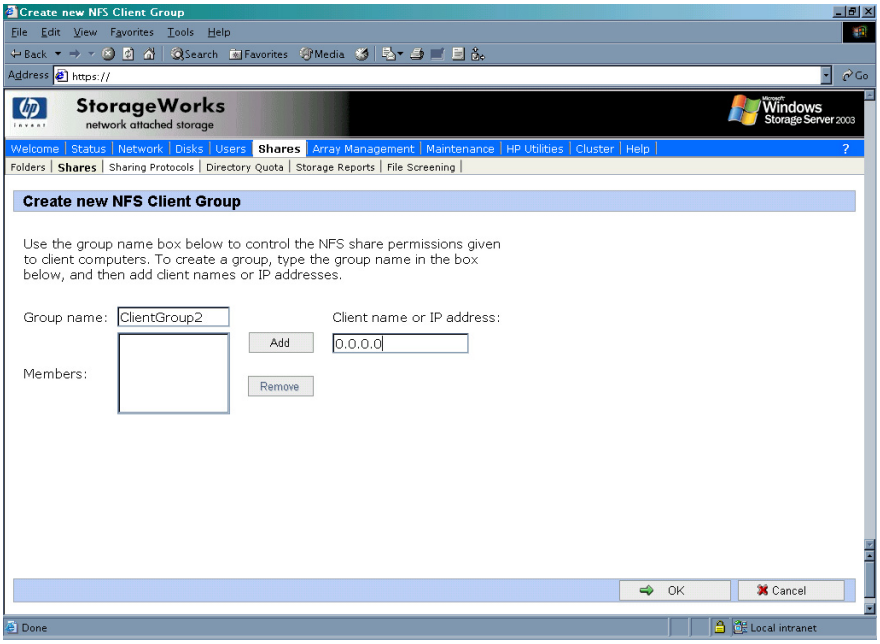

Figure 98: New NFS Client Group dialog box

2. Enter the name of the new group.
3. Enter the client name or their IP address.
4. Click **Add**. The system adds the client to the displayed list of members.
5. To remove a client from the group, select the client from the **Members** box and then click **Remove**.
6. After all clients have been added to the group, click **OK**. The **NFS Client Groups** dialog box is displayed again.

Deleting a Client Group

To delete a group:

1. From the **NFS Client Groups** dialog box, select the group to delete and click **Delete**.
2. A verification screen is displayed. Confirm that this is the correct group and then click **OK**.

   The **NFS Client Groups** dialog box is displayed again.

Editing Client Group Information

To modify the members of an existing client group:

1. From the **NFS Client Groups** dialog box, select the group to modify, and click **Edit.**

   The **Edit NFS Client Group** dialog box is displayed. Current members of the group are listed in the **Members** box.

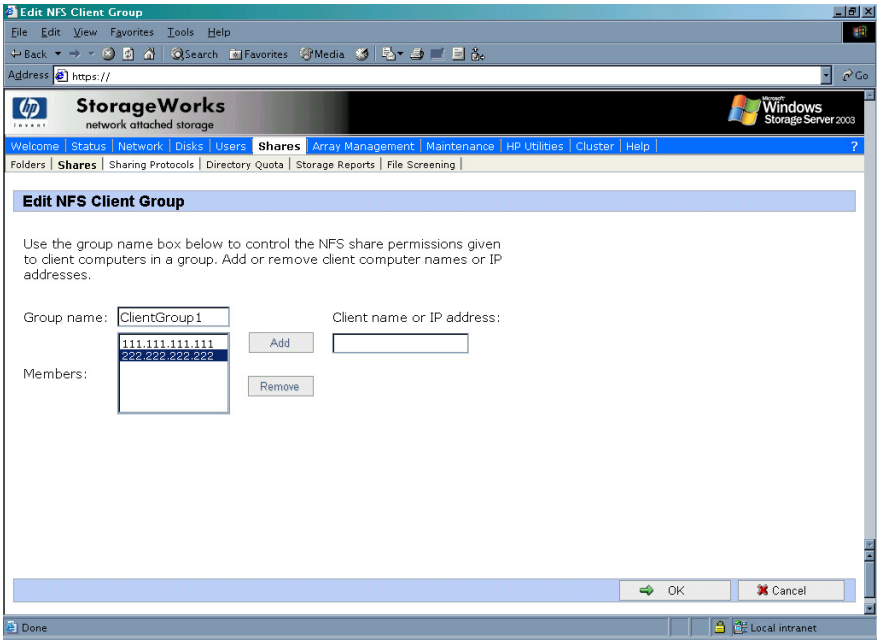

Figure 99: Edit NFS Client Groups dialog box

2. To add a client to the group, enter the client name or IP address in the **Client name** box, and then click **Add**. The client is automatically added to the **Members** list.
3. To delete a client from the group, select the client from the **Members** list, and then click **Remove**. The client is removed from the list.
4. After all additions and deletions are completed, click **OK**. The **NFS Client Groups** dialog box is displayed again.

NFS User and Group Mappings

When a fileserver exports files within a homogeneous environment, there are no problems with authentication. It is a simple matter of making a direct comparison to determine whether the user should be allowed access to the file, and what level of access to allow.

However, when a fileserver works in a heterogeneous environment, some method of translating user access is required. User mapping is the process of translating the user security rights from one environment to another.

User name mapping is the process of taking user and group identification from one environment and translating it into user identification in another environment. In the context of UNIX and NFS, user and group identification is a combination of a user ID (UID) and group ID (GID). In Windows environments, user identification is a Security ID (SID) or, in Windows Storage Server 2003, a Globally Unique Identifier (GUID).

The server grants or denies access to the export based on machine name or IP address. However, after the client machine has access to the export, user-level permissions are used to grant or deny access to user files and directories.

The NAS server is capable of operating in a heterogeneous environment, meaning that it is able to work with both UNIX and Windows clients. Because the files are stored in the native Windows NT file system, the server has to map the UNIX users to Windows users to determine the user access level of the files.

Note: User mapping is not designed to address existing user database problems in the existing environment. All UIDs and GIDs must be unique across all NIS (Network Information Service) domains and all user names must be unique across all Windows NT domains.

The NAS server supports mappings between one or more Windows domains and one or more NIS domains. The default setup supports multiple Windows NT domains to a single NIS domain. For information about users in multiple NIS domains, refer to the Supplemental Help section in the Services for NFS online help.

Types of Mappings

There are three types of mappings. These mappings are listed below in order of the most complex (with the greatest level of security) to the least complex (easiest to manage, but with little security):

- Explicit mappings
- Simple mappings
- Squashed mappings

Explicit Mappings

Explicit mappings are created by the administrator to link Windows and UNIX users. They override simple mappings and are used to map users on the different systems that have unique names.

Simple Mappings

Simple mapping is a direct comparison of user names on the Windows system and the UNIX system. If the names match, the user is assumed to be authentic, and appropriate share access is granted. Simple mapping is an option that the administrator must turn on if it is to be used.

Squashed Mappings

If the NFS server does not have a corresponding UID or GID or if the administrator has set other conditions to filter out the user, a process called squashing takes effect. Squashing is the conversion of an unmapped or filtered user to an anonymous user. This anonymous user has very restricted permissions on the system. Squashing helps administrators manage access to their exports by allowing them to restrict access to certain individuals or groups and to squash all others down to restricted (or no) access. Squashing enables the administrator to allow permissions instead of denying access to all the individuals who are not supposed to have access.

Figure 100 is a diagram showing an example of how the mapping server works for an `ls -al` command.

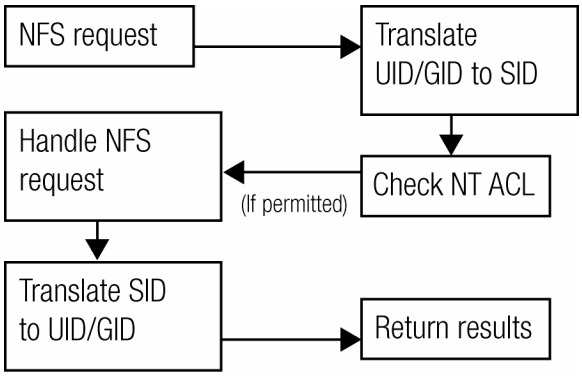

Figure 100: Mapping server `ls -al` command example

A double translation, as illustrated in Figure 100, is sometimes necessary because some commands return user ID information. For example, if the NFS request issued was an `ls -al` command, the return listing of files contains user information (the user and group that own the file). The `ls -al` command is a UNIX command. It returns a long or full listing of all files. Because this information is contained in a Windows NT Access Control List (ACL), it is not UNIX ready. The ACL information has to be converted back to UNIX UIDs and GIDs for the UNIX systems to understand and display the user information.

This second translation is not done for commands that do not return user information. For example, if the NFS request were just to read data from or write data to a file, the second translation would not be performed because there is no returning user information.

User Name Mapping Best Practices

Below is a brief list of suggested practices:

- **Back up user and group mappings**

  To avoid loss of complex advanced mappings in the case of a system failure, back up the mappings whenever the mappings have been edited or new mappings have been added.

- **Map consistently**

  Groups that are mapped to each other should contain the same users and the members of the groups should be properly mapped to each other to ensure proper file access.

  Example using User1 and Group1:

  — Make sure that the Windows User1 is mapped to the corresponding UNIX User1.

- Make sure that the Windows Group1 is mapped to the corresponding UNIX Group1.
- Make sure that User1 is a member of Group1 on both Windows and UNIX.

- **Map properly**
  - Valid UNIX users should be mapped to valid Windows users.
  - Valid UNIX groups should be mapped to valid Windows groups.
  - The mapped Windows user must have the Access this computer from the Network privilege, or the mapping will be squashed.
  - The mapped Windows user must have an active password, or the mapping will be squashed.

Creating and Managing User and Group Mappings

Note: The following sections are for a stand alone configuration. In a clustered environment, clicking **User and Group Mappings** displays a login screen for the Services for NFS Administrator.

To set up and manage user name mappings:

1. From the WebUI, select **Shares**, **Sharing Protocols**. Select **NFS Protocol** and then click **Properties**. The **NFS Properties** menu is displayed.
2. In the **NFS Properties** Menu, select **User and Group Mappings**. The **User and Group Mappings** dialog box is displayed.

   There are four tabs in the **User and Group Mappings** dialog box:

   - **General information**—Sets the mapping information source, which is either NIS or password and group files.
   - **Simple Mapping**—Indicates whether simple mappings are being used.
   - **Explicit User Mapping**—Lists exceptional user mappings that will override the simple user mappings.
   - **Explicit Group Mapping**—Lists exceptional group mappings that will override the simple group mappings.

   Each of these tabs is discussed in the following sections.
3. Enter mapping information on the appropriate tabs, then click **OK**.

General Tab

The user name mapping server translates the UNIX users into Windows users so that the server can determine user access rights to the data.

Within this initial screen, indicate whether the source of mapping information is an NIS server or is a special file with password and group information.

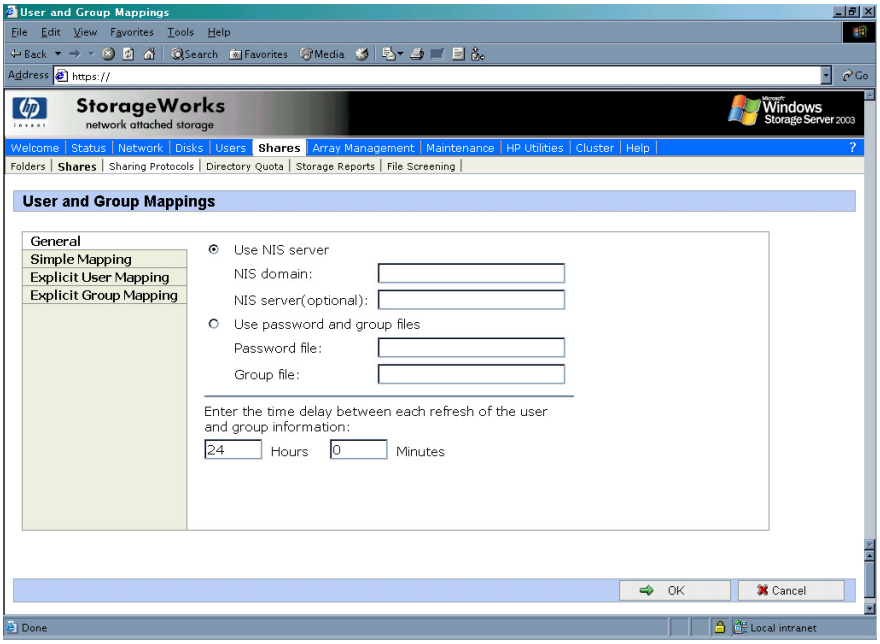

Figure 101: User and Group Mappings dialog box, General tab

From the **General** tab of the **User and Group Mappings** dialog box:

1. If an NIS server is being used:
   a. Select **Use NIS** server.
   b. Enter the NIS domain name.
   c. Enter the NIS server name. This field is optional, but recommended. In the **Hours** and **Minutes** fields, indicate how often the system will connect to the NIS domain to update the user list.
2. If custom password and group files are being used:
   a. Select **User password and group files**.
   b. Enter the path and name of the password file.
   c. Enter the path and name of the group file.
3. After this basic information is entered, click **OK**.

Simple Mapping Tab

Simple (or implicit) mapping is the first level of user name mapping. In simple mode, user and group names that match exactly in name are automatically equated.

While simple mappings are the most easily managed and are the most forthright type of map, security problems can arise. For example, if a UNIX user is coincidentally an exact match of a Windows user, the system will equate them and an inadvertent mapping will occur, granting a user inappropriate access.

- To use simple mappings, the feature must be enabled. If this feature is turned off, the administrator must manually create an explicit map for each user.
- To enable simple mapping, click the **Enable Simple Mapping** option and then select the Windows domain name.

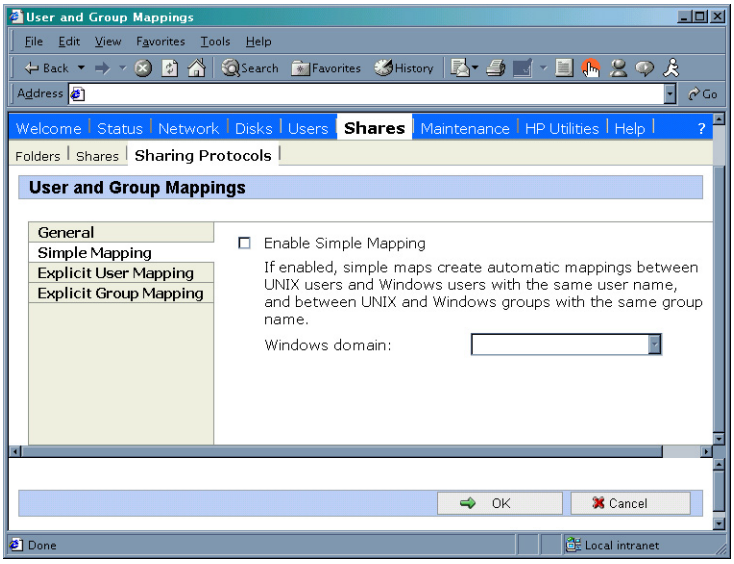

Figure 102: User and Group Mappings dialog box, Simple Mapping tab

Explicit User Mapping Tab

Explicit (or advanced) mappings allow the administrator to map any user or group manually to any other user and group. Advanced mappings override simple mappings, giving administrators the capability of using simple mapping for most users and then using advanced mappings for the users with unique names on the different systems. Alternatively, simple mapping can be disabled completely, relying solely on explicit mappings. Explicit mappings create the most secure mapping environment.

Security issues seen in simple mappings do not exist in explicit mappings. Explicit user mappings specifically correlate two users together, thus preventing the inadvertent mapping.

To enter explicit user mappings, select the **Explicit User Mapping** tab. Figure 103 is an example of the **Explicit User Mapping** tab.

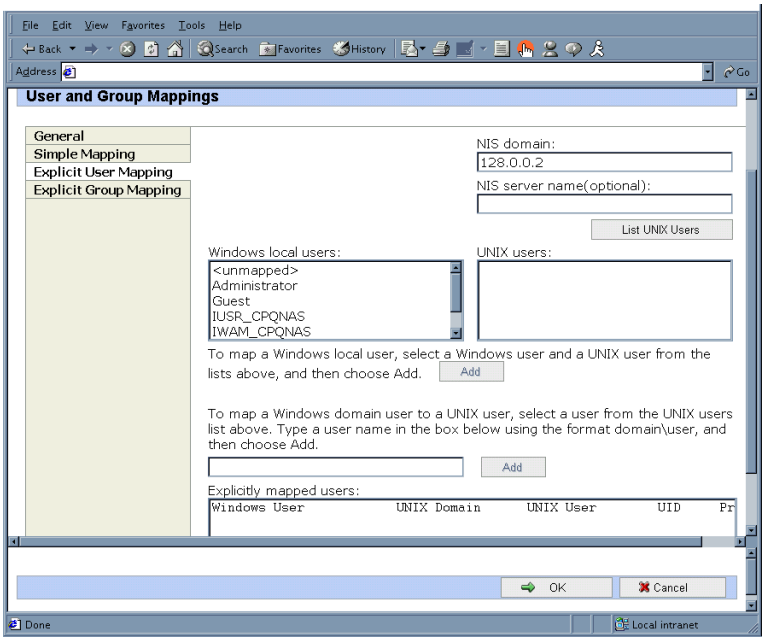

Figure 103: User and Group Mappings dialog box, Explicit User Mapping tab

To create explicit user mappings:

1. Click the **List UNIX Users** button to populate the UNIX users box.
2. To map a local Windows user to a UNIX user, highlight the **Windows user** in the Windows local users box and highlight the UNIX user that you want to map, and then click **Add**. The **Explicitly mapped users** box at the bottom of the screen is populated with the new mappings. Repeat this process until all desired users have been mapped.
3. To map a domain Windows user to a UNIX user, enter the domain and the user name in the box in the middle of the screen (use the Domain\username format) and highlight the UNIX user that you want to map, and then click **Add**. The map is added to the **Explicitly mapped users** box at the bottom of the screen. Repeat this process until all desired users have been mapped.
4. To map multiple Windows users to one UNIX user, one of the mapped Windows users must be set as the primary mapping. To indicate which user map is the primary mapping, highlight the desired map in the **Explicitly mapped users** box, and then click the **Set Primary** button.
5. To delete a map, highlight the map in the **Explicitly mapped users** box, and then click the **Remove** button.
6. After all entries are completed, click **OK** to activate the new entries.

Explicit Group Mapping Tab

To enter explicit group mappings, select the Explicit Group Mapping tab. Figure 104 is an example of the **Explicit Group Mapping** tab.

Explicit mappings allow the administrator to map any user or group manually to any other user and group. Explicit mappings override simple mappings, giving administrators the capability of using simple mapping for most groups and then using explicit mappings to make changes to simple mappings. Simple mapping can be turned off for greater security.

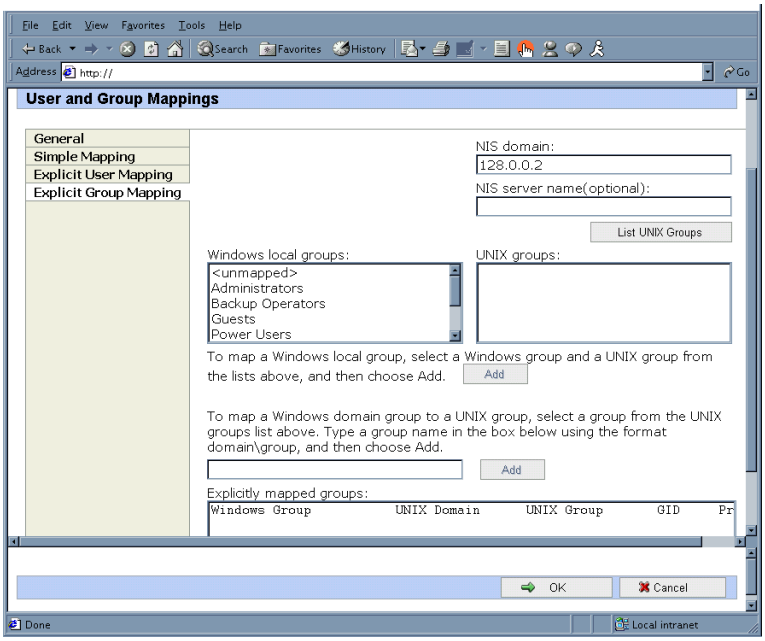

Figure 104: User and Group Mappings dialog box, Explicit Group Mapping tab

To create explicit group mappings:

1. Click the **List UNIX Groups** button to populate the **UNIX Groups** box.
2. To map a local Windows group to a UNIX group, highlight the Windows group in the Windows local groups box and highlight the UNIX group to map, and then click **Add**. The **Explicitly mapped groups** box at the bottom of the screen is populated with the new mappings. Repeat this process until all desired groups have been mapped.
3. To map a domain Windows group to a UNIX group, enter the domain and the group name in the box in the middle of the screen (use the Domain\groupname format) and highlight the UNIX group to map, and then click **Add**. The map is added to the **Explicitly mapped groups** box at the bottom of the screen. Repeat this process until all desired groups have been mapped.
4. To map multiple Windows groups to one UNIX group, one of the Windows groups must be set as the primary mapping. Therefore, to indicate which group map is the primary mapping, highlight the desired map in the **Explicitly mapped groups** box, and then click the **Set Primary** button.
5. To delete a map, highlight the map in the **Explicitly mapped groups** box and then click the **Remove** button.
6. After all entries are completed, click **OK** to activate the new entries.

Backing up and Restoring Mappings

The user name-mapping server has the capability to save and retrieve mappings from files. This capability is useful for backing up mapping settings prior to making changes and for exporting the mapping file from one server to others, using the same mapping information.

The user name-mapping server can save existing mappings to a file or load them from a file and populate the mapping server. This feature is found in the **Map Maintenance** tab of the **User Name Mapping** screen, as shown in Figure 105.

Use **Remote Desktop** to access the **NAS Management Console**, click **File Sharing**, **Microsoft Services for Network File System**. Click **User Name Mapping**, then **Map Maintenance**.

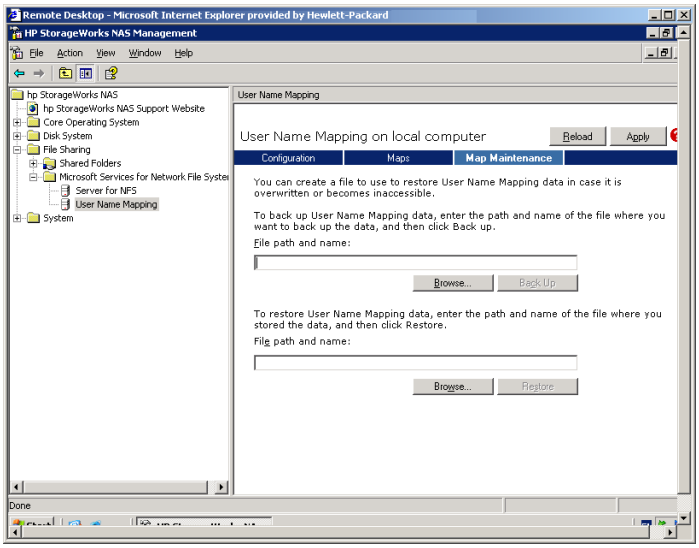

Figure 105: User Name Mapping screen, Map Maintenance tab

Backing up User Mappings

1. Select the **Map Maintenance** tab from the **User Name Mapping** screen.
2. Type the path and name of the file to be used for backup in the File path and name field or click **Browse** to locate the file.

Note: If the file is being created for the first time, follow these steps:

---

1. Browse to the target directory.
2. Right-click in the file listing pane, select **New**, **Text Document**. Enter a name for the file and then press **Enter**.
3. Double-click the new file to select it.
4. Click **Backup**.

Restoring User Mappings

User mappings can be restored using the following procedures.

1. Select the **Map Maintenance** tab from the **User Name Mapping** screen.

2. Type the path and name of the file in the File path and name field or click **Browse** to locate the file.
3. After locating the file, click **Restore**.

Creating a Sample NFS File Share

HP recommends performing the following tests to verify that the setup of the shares, user mappings, and permissions grant the desired access to the NFS shares.

1. Create an NFS share. NFS Shares are All Machines, read-only by default.

   See "NFS File Shares" earlier in this chapter for information on creating shares.
2. Create NFS client groups if desired. See "NFS Client Groups" earlier in this chapter.
3. Verify that the NFS share exists.

   Use Remote Desktop to log into the NAS server and access the command line interface:

   `nfsshare <sharename>` (sharename represents the name of the share)
4. Map a user. When creating Active Directory/Domain mappings, ensure that the NFS Authentication software is installed on the domain controllers that have user name mappings. See "Installing NFS Authentication Software on the Domain Controllers and Active Directory Domain Controllers" section. Also, see "User and Group Mappings" in this chapter for instructions on setting up user name mappings.

   When planning to allow only anonymous access to an NFS share, setting up user name mappings is not necessary. See the section, "Anonymous Access to an NFS Share" in this chapter for additional information.
5. Verify the NTFS permissions are correct on the NFS share. If the NFS share was assigned All Machines read write, then the NTFS ACLs on the NFS share must allow read/write permissions for the user or group.

   Example: *e:\share1* is the name of the NFS share and share1 has All Machines read write permissions. Verify that the NTFS permissions on e:\share1 are List Folder/Read Data, Create File/Write Data, Create Folders/Append Data, Write Attributes, and Delete Subfolders and Files. This can be verified by opening up Windows Explorer on the NAS Desktop and right-clicking e:\share1 then clicking **Properties**. Next, click the **Security** tab. Then click **Advanced**. Highlight the user or group that permissions are being assigned to then click **Edit**. There will be check boxes next to the NTFS permissions that are assigned. Make sure mapped users and groups correlate to the users or groups that have the NTFS permissions assigned. See the section "Understanding NTFS and UNIX Permissions" in this chapter for more information.
6. Verify that the mappings exist.

   Use Remote Desktop to log in to the NAS server and access the command line interface:

   `mapadmin list -all`
7. On the Linux/UNIX system, use the mapped user to create a file.
   a. As the root user, mount the share:

      `mount -t nfs <nfs server IP address:/nfs share> /mount point`
   b. Log in as a mapped user.
   c. Change directories to the mount-point directory.

d. Create the file as the mapped user (example: *file1*).

8. Verify that the same permissions are set up for the user on both the UNIX side and the Windows side.
   a. List the permissions on the UNIX side:

      ```
      ls -l /mount-point/file1
      ```

      (Example screen display: `-r--r----- unixuser1 unixgroup1`)

   b. List the permissions on the Windows side: (change to the *nfs* share directory)

      From a command line interface accessed from Remote Desktop on the NAS server:

      ```
      cacls file1
      ```

      (Example display: `DOMAIN1\Windowsuser1:R`)

   c. Compare and verify the permissions from UNIX and Windows.

Remote Desktop

In addition to the WebUI, Remote Desktop is available for remote administration of Services for UNIX. This service let users connect to machines, log on, and obtain command prompts remotely. See Table 12 for a list of commonly used commands.

Using Remote Desktop

Microsoft Remote Desktop can be used to remotely access the NAS server desktop. This provides the administrator flexibility to automate setups and other tasks. Services for NFS file-exporting tasks and other Services for NFS administrative tasks can be accomplished using Remote Desktop to access the Services for NFS user interface from the NAS Desktop or from a command prompt.

Remote Desktop is included in the WebUI of the NAS server. To open a Remote Desktop session, from the WebUI, select **Maintenance**, **Remote Desktop**. See the "Remote Access Methods and Monitoring" chapter for information on setting up and using Remote Desktop.

Table 12 describes some common Services for NFS commands.

Table 12: Command Line Interface Command Prompts

Command	Function
`nfsstat /?`	Learn about viewing statistics by NFS operation type
`showmount /?`	View the format of the command to display NFS export settings on NFS servers
`showmount -a`	View users who are connected and what they currently have mounted
`showmount -e`	View exports from the server and their export permissions
`rpcinfo /?`	Learn how to display Remote Procedure Call (RPC) settings and statistics
`mapadmin /?`	View how to add, delete, or change user name mappings
`nfsshare /?`	Learn how to display, add, and remove exported shares

NetWare File System Management

9

File and Print Services for NetWare (FPNW) is one part of the Microsoft software package called Services for NetWare. The most common use of the NetWare network operating system is as a file and print server. Customers using NetWare as the platform to host their file and print services have become accustomed to its interface from both a user and an administrator point of view and have built up an investment in NetWare file and print services. File and Print Services for NetWare helps customers preserve their NetWare skill set while consolidating the number of platforms. This reduces hardware costs and simplifies file and print server administration by making the NAS server emulate a NetWare file and print server. FPNW eases the addition of the NAS server into a mixed infrastructure by providing a NetWare user interface (UI) to a Windows Storage Server 2003-based server; administrators and users see their same, familiar NetWare UI. Additionally, the same logon for clients is maintained without a need for any client configuration changes.

This service also provides the ability to create Novell volumes, which are actually NTFS shares, from which users can map drives and access resources. Novell Login scripts are supported on the NAS server or through an existing NDS (Novell Directory Services) account.

Note: NetWare is not a clusterable protocol. With NetWare on both nodes of the cluster, the shares will not failover as the protocol is not cluster-aware.

Note: IPX/SPX protocol is required on the Novell servers.

Topics discussed in this chapter include:

- Installing Services for NetWare
- Managing File and Print Services for NetWare
- Creating and Managing NetWare Users
- Managing NCP Volumes (Shares)

Installing Services for NetWare

The installation of FPNW on the NAS server allows for a smooth integration with existing Novell servers. FPNW allows a Windows Storage Server 2003-based server to emulate a NetWare file and print server to users, clients, and administrators. This emulation allows authentication from Novell clients, the use of Novel logon scripts, the creation of Novell volumes (shares), the use of Novell file attributes, and many other Novell features.

Additional information on Microsoft Directory Synchronization Services and the File Migration Utility can be found at:

www.microsoft.com/WINDOWS2000/guide/server/solutions/NetWare.asp

---

Note: The printing capabilities of File and Print Services for NetWare are not supported on the NAS server.

To install Services for NetWare:

1. From the desktop of the NAS server, click **Start > Settings > Network Connections > Local Area Connection**, and then click **Properties**.
2. Click **Install**. The **Select Network Component Type** dialog box is displayed.

   Figure 106 is an example of the **Select Network Component Type** dialog box.

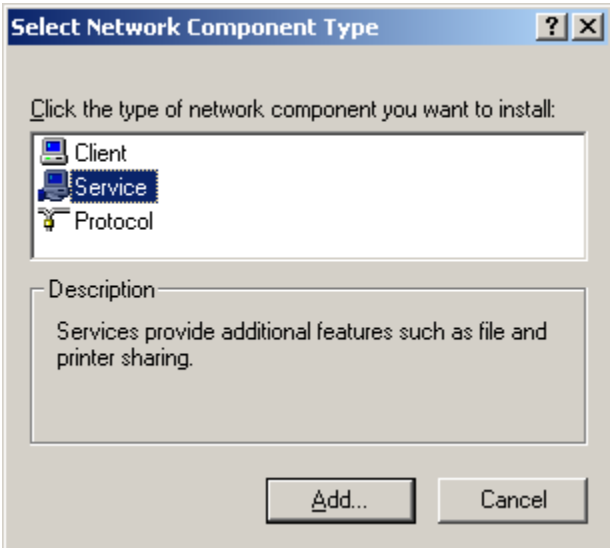

Figure 106: Local Area Connection Properties page, Install option

3. Select **Service** and click **Add**.
4. Click the **Have Disk** icon and navigate to the location of **Services for NetWare**.

   Services for NetWare is located in the path:
   c:\hpnas\components/SFN5.02/fpnw/netsfn.inf.

   Click **Open**.
5. Click **OK**.

   File and Print Services for NetWare should now be displayed as an option to install.

6. Select **File and Print Services for NetWare** and click **OK**.

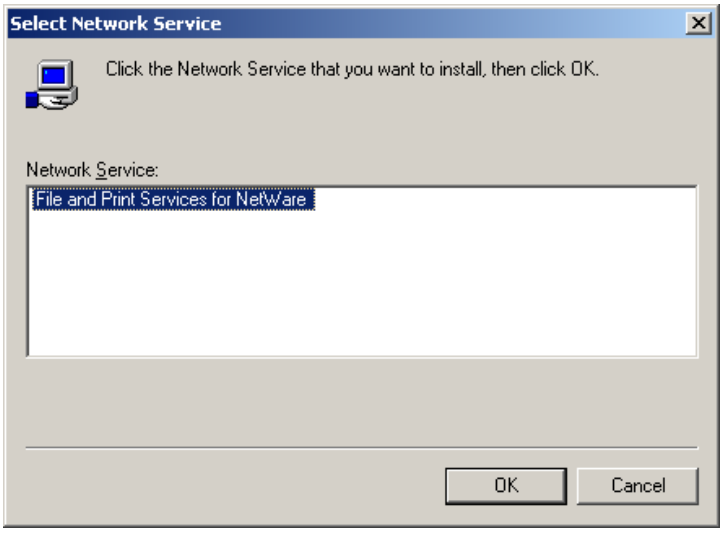

Figure 107: Select network service

7. Click **OK**.

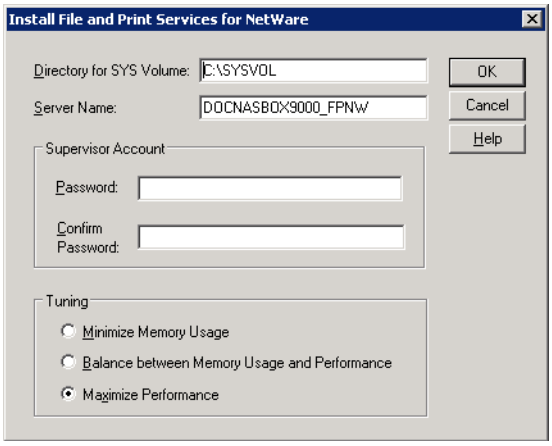

Figure 108: Install File and Print Services for NetWare

8. Select or change items as necessary.
9. Click **OK**.

The system must be rebooted for the changes to take effect.

Managing File and Print Services for NetWare

To access FPNW:

1. From the desktop of the NAS server, click **Start > Settings > Control Panel > Administrative Tools > Server Manager**.
2. Select **FPNW**, then **Properties**.

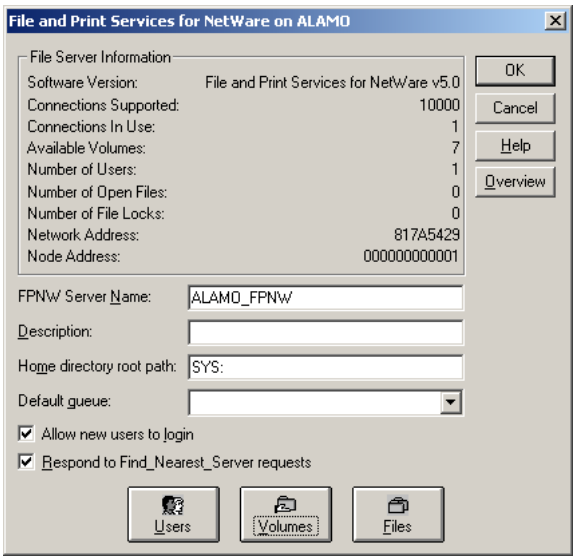

Figure 109: File and Print Services for NetWare screen

3. Enter an **FPNW Server Name** and **Description**.

   This name must be different from the server name used by Windows or LAN Manager-based clients to refer to the server. If you are changing an existing name, the new name will not be effective until you stop and restart **File and Print Services for NetWare**. For example, in Figure 109 the Windows server name is Alamo and the FPNW server name is Alamo_FPNW.

4. Indicate a **Home directory root path**.

   This path is relative to where the Sysvol volume has been installed. This will be the root location for the individual home directories. If the directory specified does not already exist, it must first be created.

5. Click **Users** to:

   See connected users, disconnect users, send broadcast messages to all users connected to the server, and to send a message to a specific user.

6. Click **Volumes** to:

   See users connected to specific volume and to disconnect users from a specific volume.

7. Click **Files** to:

   View open files and close open files.

Creating and Managing NetWare Users

To use Services for NetWare, the Novell clients must be entered as local users on the NAS server.

Adding Local NetWare Users

1. From the NAS server desktop, click the **NAS Management Console** icon, click **Core Operating System**, and then click **Local Users and Groups**.
2. Right-click the **Users** folder and then click **New User**.

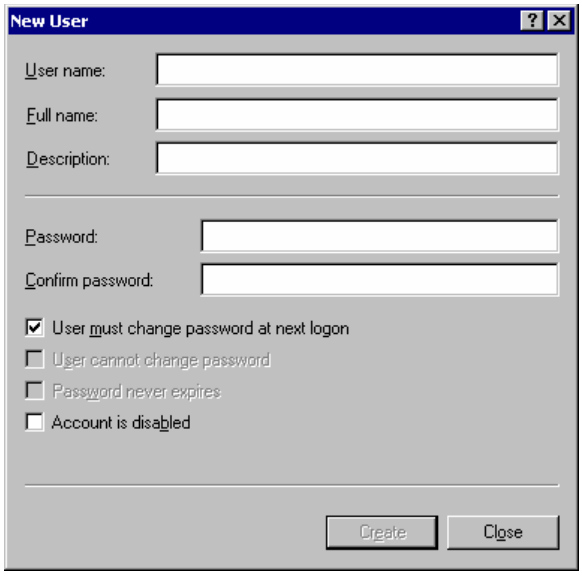

Figure 110: New User dialog box

3. Enter the user information, including the user's User name, Full name, Description, and Password. Click **Create**.
4. Repeat these steps until all NetWare users have been entered.

Enabling Local NetWare User Accounts

1. In the **Users** folder (NMC, Core Operating System, Local Users and Groups), right-click an NCP client listed in the right pane of the screen and then click **Properties**.
2. Select the **NetWare Services** tab.

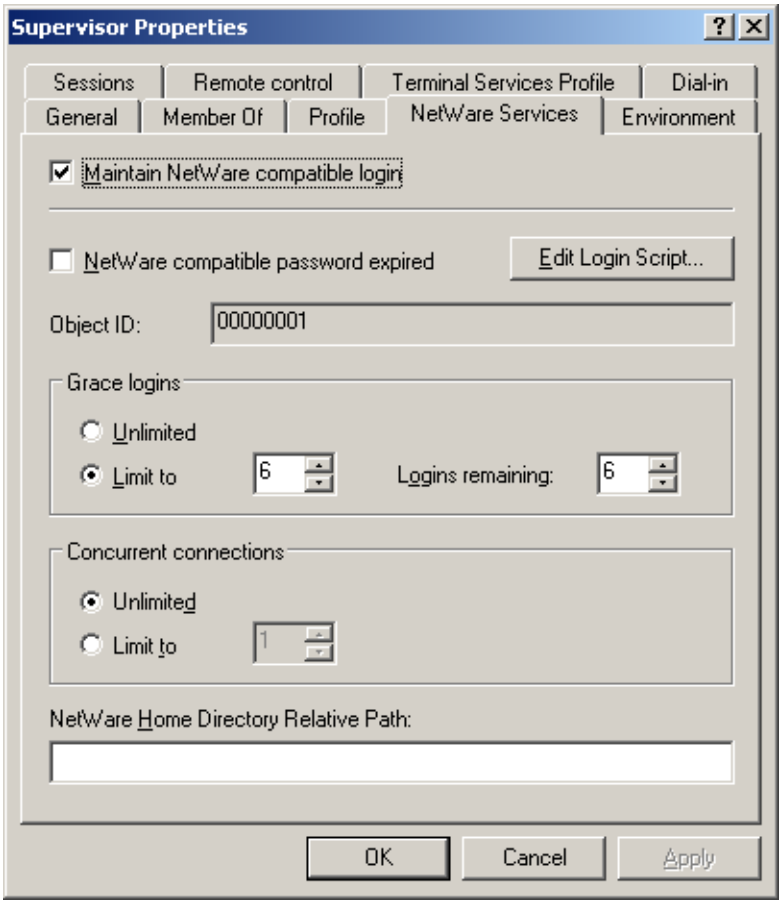

Figure 111: NetWare Services tab

3. Select **Maintain NetWare compatible login**.
4. Set other NetWare options for the user and click **OK**.

Note: The installation of File and Print Services for NetWare will also create a supervisor account, which is used to manage FPNW. The supervisor account is required if the NAS server was added as a bindery object into NDS.

Managing NCP Volumes (Shares)

NCP file shares are created in the same manner as other file shares; however, there are some unique settings. NCP shares can be created and managed using the NAS Management Console.

Note: NCP shares can be created only after Microsoft Services for NetWare is installed. See the previous section "Installing Services for NetWare" for instructions on installing SFN.

Creating a New NCP Share

To create a new file share:

1. From the NAS server desktop, choose **Start > Settings > Control Panel > Administrative Tools > Server Manager**.
2. Choose **FPNW > Shared Volumes**.
3. Click **Create Volume**.

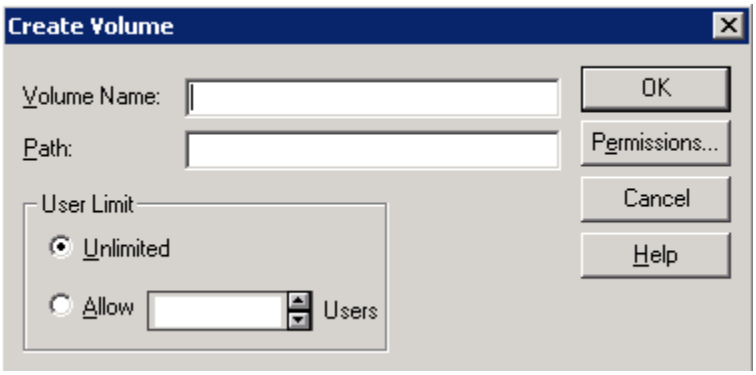

Figure 112: Create Shared Folder dialog box

4. Specify the volume name and path.
5. Click **Permissions** to set permissions.

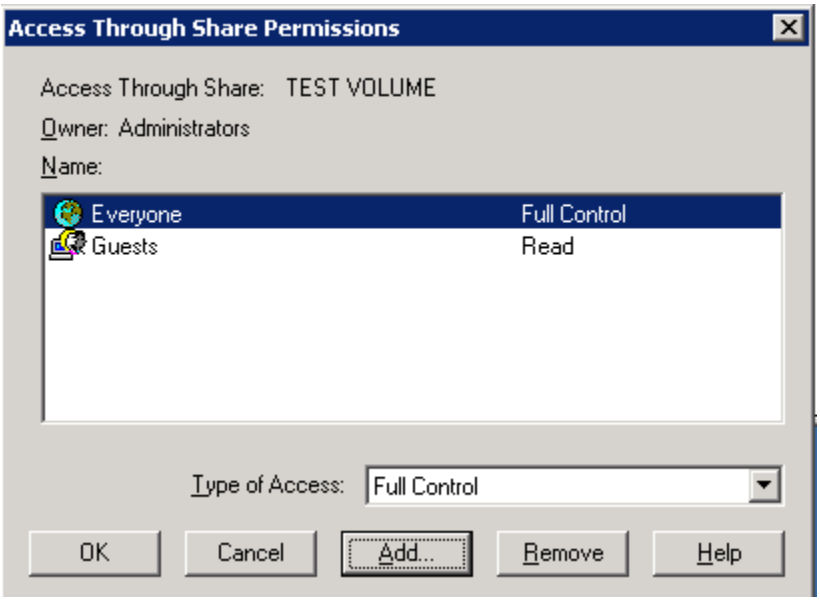

Figure 113: Share permissions dialog box

6. Click **Add** to add additional users and groups, and to set their permissions.

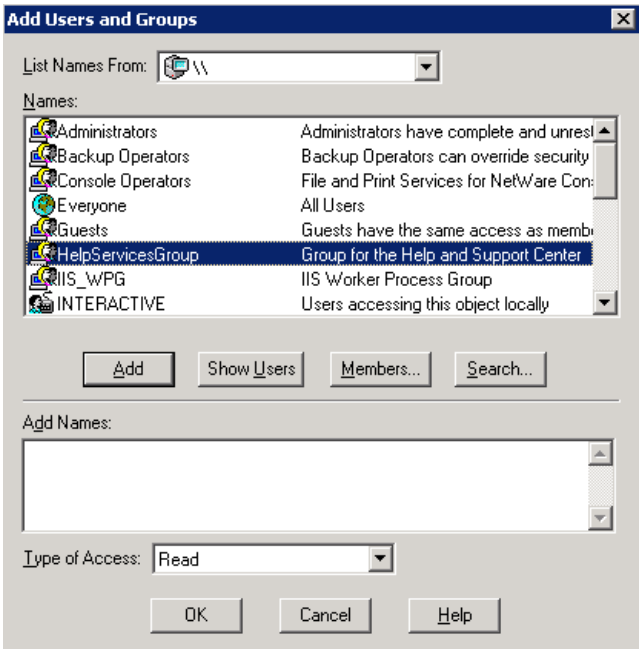

Figure 114: Add Users and Groups dialog box

7. Highlight the desired user or group, then click **Add**.
8. Select the Type of Access from the drop down list.

Note: Type of Access can also be set from the Access Through Share Permissions dialog box.

---

9. Click **OK** when all users and groups have been added.
10. Click **OK** on the Create Volume dialog box.
11. Click **Close**.

Modifying NCP Share Properties

To modify a file share:

1. From the NAS server desktop, choose **Start > Settings > Control Panel > Administrative Tools > Server Manager**.
2. Choose **FPNW > Shared Volumes**.
3. Highlight the volume to modify.
4. Click **Properties**.

Cluster Administration

10

One important feature of the HP StorageWorks NAS server is that it can operate as a single node or as a cluster. This chapter discusses cluster installation and cluster management issues. Some of these topics are discussed or mentioned elsewhere in this guide. The discussion in this chapter is more detailed than other references and addresses the unique administration procedures for operating in a clustered environment.

Cluster Overview

As introduced in chapter 1, two server heads (nodes) can be connected to each other and deployed as a no single point of failure (NSPOF) dual redundant cluster. The nodes are connected by a crossover cable and are each connected to network switches or hubs. This connection allows communication between the nodes to track the state of each cluster node. Each node sends out periodic messages to the other node; these messages are called heartbeats. If a node stops sending messages, the cluster service will fail over any resources that the node owns to the other node. For example, if the node that owns the Quorum disk is shut down for any reason, its heartbeat will stop. The other node detects the lack of the heartbeat and takes over ownership of the Quorum disk and the cluster.

Multi Node Support Beyond Two Nodes

The NAS 4000s and 9000s devices may be deployed in multi node clustering beyond two nodes. Refer to the associated Storage Array documentation to determine the number of nodes supported by the array under Windows Storage Server 2003. While the discussion presented in this guide addresses only two nodes, additional nodes may be added into the cluster. Considerations for additional fiber path connections and the private network should be made. In the case of the private network, a hub or switch is required since the cross over cable is no longer applicable.

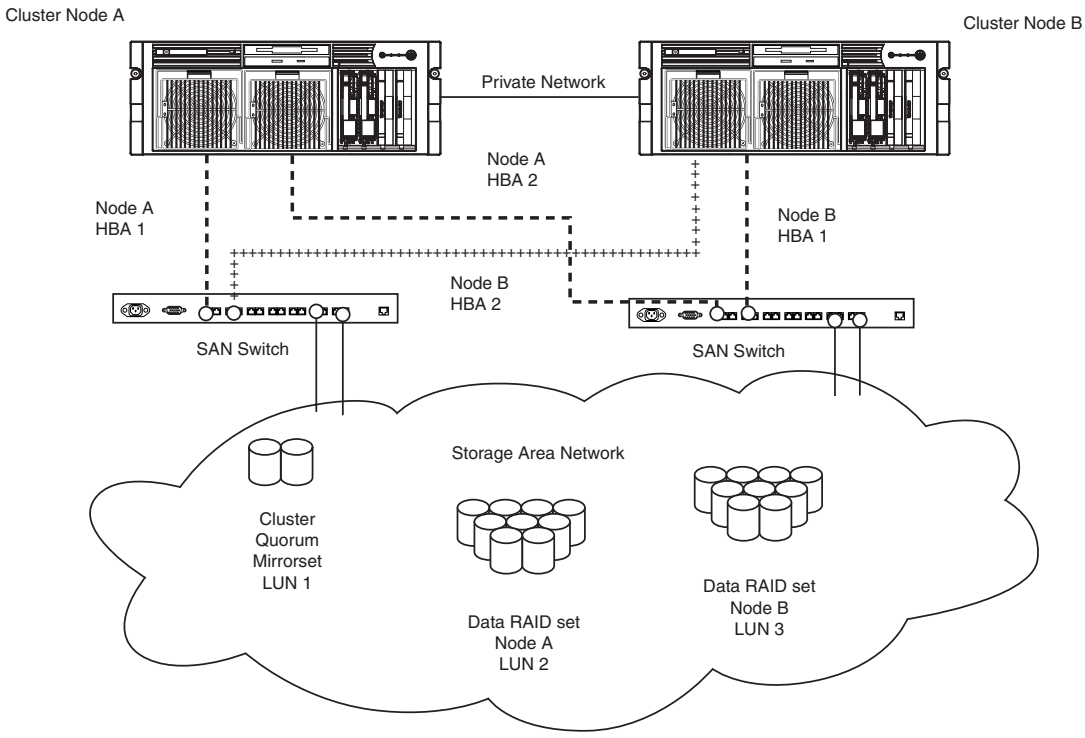

Figure 115: NAS server cluster diagram

Cluster Terms and Components

This section provides brief definitions of clustering terms. This information provides basic knowledge of clusters and the terminology used throughout this document.

Nodes

The most basic parts of a cluster are the server heads. A server node is any individual computer in a cluster or a member of the cluster. If the NAS device is a member of a cluster, then the server heads are referred to as nodes.

Resources

Hardware and software components that are managed by the cluster service are called cluster resources. Cluster resources have three defining characteristics:

- They can be brought online and taken offline.
- They can be managed in a cluster.
- They can be owned by only one node at a time.

Examples of cluster resources are IP addresses, network names, physical disk resources, and file shares.

Virtual Servers

A virtual server is a cluster group that consists of a static IP Address resource and a Network Name resource. Several virtual servers can be created. By assigning ownership of the virtual servers to the different server nodes, the processing load on the NAS devices can be distributed between the nodes of a cluster.

The creation of a virtual server allows resources dependant on the virtual server to fail over and fail back between the cluster nodes. File Share and physical disks resources are assigned to the virtual server to ensure non disruptive service of file shares to the clients.

Failover

Failover of cluster groups and resources happens:

- when a node hosting the group becomes inactive. A shutdown of cluster service or a loss of power can cause a failover.
- when all of the resources within the group are dependent on one resource and that resource fails.
- when an administrator forces a failover.

A resource and all of its dependencies must be located in the same group so that if a resource fails over, all of its dependent resources fail over.

When a resource is failed over, the cluster service performs certain procedures. First, all of the resources are taken offline in an order defined by the resource dependencies. Secondly, the cluster service attempts to transfer the group to the next node on the preferred owners list. If the transfer is successful, the resources are brought online in accordance with the resource dependency structure.

The system failover policy defines how the cluster detects and responds to the failure of individual resources in the group. After a failover occurs and the cluster is brought back to its original state, failback can occur automatically based on the policy. After a previously failed node comes online, the cluster service can fail back the groups to the original host. The failback policy must be set before the failover occurs so that failback works as intended.

Quorum Disk

Each cluster must have a shared disk called the Quorum disk. This physical disk in the common cluster disk array plays a critical role in cluster operations. The Quorum disk offers a means of persistent storage. The disk must provide physical storage that can be accessed by any node in the cluster. If a node has control of the quorum resource upon startup, it can initiate the cluster. In addition, if the node can communicate with the node that owns the quorum resource, it can join or remain in the cluster.

The Quorum disk maintains data integrity by:

- storing the most current version of the cluster database.
- guaranteeing that only one set of active communicating nodes is allowed to operate as a cluster.

Cluster Concepts

Microsoft cluster concepts are rather straight forward when explained through a diagram. Figure 116 illustrates a typical cluster configuration with the corresponding storage elements. The diagram progresses from the physical disks to the file shares, showing the relationship between both the cluster elements and the physical devices underlying them. While the diagram only illustrates two nodes, the same concepts apply for multi node deployments.

Sequence of Events for Cluster Resources

The sequence of events in the diagram includes:

1. Physical disks are combined into RAID arrays and LUNs.
2. LUNS are designated as basic disks, formatted and assigned a drive letter via Disk Manager
3. Physical Disk resource are created for each basic disk inside cluster administrator.
4. Directories and folders are created on assigned drives.
5. Cluster components (virtual servers, file shares) are created, organized in groups, and placed within the folders using cluster administrator exclusively.

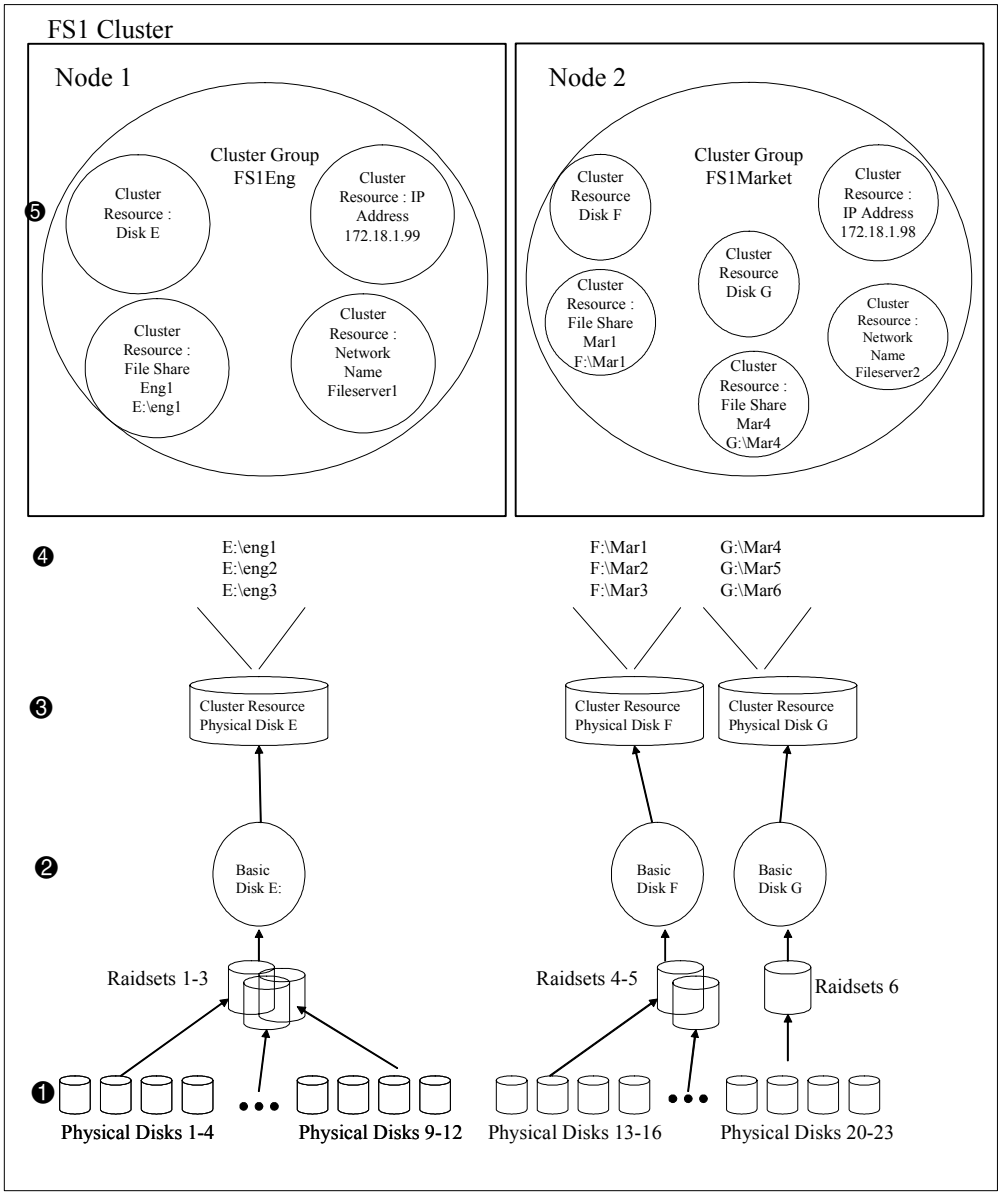

Figure 116: Cluster concepts diagram

Hierarchy of Cluster Resource Components

The cluster components are referred to as resources and are placed together in groups. Groups are the basic unit of failover between nodes. Resources do not failover individually, rather they failover with the group in which they are contained.

In Figure 116 it is depicted as follows:

- Physical Disk resources are placed in a group and relate to the basic disk created through the WebUI. It should be noted that when a Physical Disk resource is created through Cluster Administrator a corresponding group should be created for the resource to reside in. Groups are the basic unit of failover on a cluster.
- File Share resources are placed in a group and relate to the actual directory on the drive on which the share is being created.
- An IP Address resource is formed in the group and relates to the IP address by which the group's virtual server is identified on the network.
- A Network Name resource is formed in the group and relates to the name published on the network by which the group is identified.
- A Virtual Server is a group containing an IP Address resource and a Network Name resource. File share and disk resources assigned to this virtual server group can transition from one node to the other during failover conditions.
- The Group is owned by one of the nodes of the cluster, but may transition to the other nodes during failover conditions.

The diagram illustrates a cluster containing two nodes. Each node has ownership of one group. Contained within each group are singular file shares that are known on the network by the associated Network Name and IP address. In the specific case of Node1, file share Eng1 relates to *E:\Eng1*. This file share is known on the network as *\\Fileserver1\Eng1* with an IP address of 172.18.1.99. *E:\Eng1* relates to the actual Basic Disk E: containing a directory *Eng1*.

For cluster resources to function properly, two very important requirements should be adhered to:

- Dependencies between resources of a group must be established. Dependencies determine the order of startup when a group comes online. In the above case, the following order should be maintained:

  1. File Share—dependent on Physical Disk Resource
  2. NFS File Share—dependent on Physical Disk Resource and Network Name
  3. Network Name—dependent on IP Address

  Failure to indicate the dependencies of a resource properly may result in the file share attempting to come online prior to the physical disk resource being available, resulting in a failed file share.
- Groups should have a Network Name resource and an IP Address resource. These resources are used by the network to give each group a virtual name. Without this virtual reference to the group, the only way to address a share that is created as a clustered resource is by node name. Physical node names do not transition during a failover, whereas virtual names do.

For example, if from a client a network share map F: was established and assigned to *Node1**Eng1* instead of *Fileserver1**Eng1*, when Node1 fails and Node2 assumes ownership, the map will become invalid because the reference in the map is to *Node1*. If the map were created to the virtual name and Node1 were to fail, the map would still exist when the group associated with Eng1 failed over to Node2.

The previous diagram is an example and is not intended to imply limitations of a single group or node. Groups can contain multiple physical disks resources and file shares and nodes can have multiple groups, as shown by the group owned by Node2.

Cluster Planning

Clustering the NAS 4000s or 9000s greatly enhances the availability of file service by enabling file shares to fail over to additional NAS devices, if problems arise. Clients see only a brief interruption of service as the file share resource transitions from one server node to the other.

Requirements for taking advantage of clustering include:

- Storage planning
- Network planning
- Protocol planning

Storage Planning

For clustering, a storage unit (LUN) must be designated for the cluster and configured as a mirrorset. This LUN is used for the Quorum disk. The Quorum disk is the shared storage used by the cluster nodes to coordinate the internal cluster state.

One or more RAID arrays are dedicated to each cluster node for data storage. Each cluster node will assume ownership of at least one physical disk resource. That owner node will serve all shares within that physical disks resource, until a failover condition occurs. When a failover occurs, the physical disk resource and all associated shares will transition over to the remaining nodes and will remain there until the other node is returned to service. Some types of shares are not cluster aware and will not be available during a failover condition. See the "Protocol Planning" section for additional information.

To prepare a basic disk for use in a cluster, a cluster group for each basic disk should be created to allow each resource to failover separately. Once the group is created, a physical disk resource is created in each of the groups. Cluster groups may contain more than one physical disk depending on the site-specific requirements. This physical disk resource is required for basic disk to successfully work in a cluster environment protecting it from simultaneous access from each node.

Note: The LUN underlying the basic disk should be presented to only one node of the cluster using selective storage presentation SAN switch zoning or having only one node online at all times until such times as the physical resource for the basic disk is established.

In preparing for the cluster installation:

- All software components listed in the SAN connection tool must be installed and the fiber cables attached to the HBA(s) before the cluster installation is started.
- All shared disks, including the quorum disk, must be accessible from both nodes. When testing connectivity between server and LUN, only one server should be given access to the LUN at a time or the non-testing server should be powered off.
- All shared disks must be configured as basic (not dynamic).
- All partitions on the disks must be formatted as NTFS.

Network Planning

Clusters require more sophisticated networking arrangements than a stand alone NAS device. For example, because a cluster must be deployed into a domain environment, workgroups are not supported. A Windows NT domain or Active Directory domain must be in place to contain the cluster names, virtual server names, and user and group information. A cluster cannot be deployed into a non domain environment.

All cluster deployments have at least seven network addresses and network names:

- The cluster name (Unique NETBIOS Name) and IP address
- Node A's name and IP address
- Node B's name and IP address
- At least one virtual server name and IP address for Node A
- At least one virtual server name and IP address for Node B
- Cluster Interconnect static IP addresses for Node A and Node B

In multi node deployments additional network addresses are required. For each additional node, three static IP addresses are required.

Virtual names and addresses are the only identification used by clients on the network. Because the names and addresses are virtual, their ownership can transition from one node to the other during a failover, preserving access to the shares on the virtual disks.

In addition, a cluster will use at least two network connections on each node:

- The cluster interconnect or “heartbeat” crossover cable connects to the first network port on each cluster node. In more than two node deployments, a private VLAN on a switch or hub is required for the cluster interconnect.
- The client network subnet connects to a second network port on each cluster node. The cluster node names and virtual server names will have IP addresses residing on these subnets.

Note: If the share is to remain available during a failover, each cluster node must be connected to the same network subnet. It is impossible for a cluster node to serve the data to a network to which it is not connected.

Protocol Planning

The NAS 4000s and 9000s both support many file sharing protocols, including sharing protocols for Windows, UNIX, Linux, Novell, Macintosh, Web, and FTP clients. However, not all of these protocols can take advantage of clustering. If a protocol does not support clustering, the share will not be available to the clients until the owner cluster node is brought back online.

HP recommends placing cluster aware and non cluster aware protocols on different file shares.

Use the information in Table 13 to determine whether it is advantageous to use clustering.

Table 13: Sharing Protocol Cluster Support

Protocol	Client Variant	Cluster Aware (supports failover)	Supported
CIFS/SMB	Windows NT Windows 2000 Windows 95 Windows 98 Windows ME	Yes	Yes
NFS	UNIX Linux	Yes	Yes
HTTP	Web	No	Yes
FTP	Many	Yes	Yes
NCP	Novell	No	Yes
AppleTalk	Apple	No	No

Note: AppleTalk is not supported on clustered disk resources. AppleTalk requires local memory for volume indexing. On failover events, the memory map is lost and data corruption can occur.

Preparing for Cluster Installation

This section provides the steps necessary to cluster HP StorageWorks NAS servers.

Before Beginning Installation

Confirm that the following specifications have been met before proceeding:

- The SAN connection tool must be completed and all the necessary software components for connecting to the desired storage must be installed before the configuration of cluster services.
- It is required that at least one LUN has been presented for the configuration of the Quorum disk. This LUN must be created from shared storage and must be at least 50 MB. (500 MB is recommended). Additional LUNS may also be presented for use as shared disk resources.
- Cluster configurations should be deployed with dual data paths for high availability. Dual data paths from each node enable a path failure to occur that does not force the failover of the node. Clusters can be configured with single path, but if a failure in the path does occur, all of the node resources will be failed to the non-affected node.

HP StorageWorks NAS Software Updates

Caution: The version of Storage Manager (Quota Management software located under the Shares tab) that comes installed on the server is not supported in a cluster. It must be uninstalled prior to creating the cluster.The uninstall tool is located in the Cluster Installation Guide under the Cluster tab in the WebUI. See Figure 117.

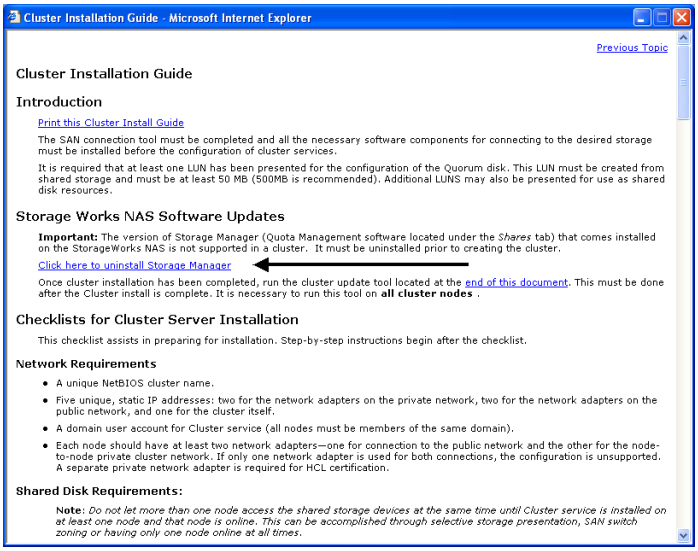

Figure 117: Uninstall Storage Manager

After cluster installation has been completed, run the cluster update tool located in the Cluster Installation Guide in the WebUI. The Cluster Installation Guide is located under the Cluster tab. This must be done after the cluster installation is complete. It is necessary to run this tool on all cluster nodes.

Checklists for Cluster Server Installation

These checklists assist in preparing for installation. Step-by-step instructions begin after the checklists.

Network Requirements

- A unique NetBIOS cluster name.
- For each node deployed in the cluster the following static IP addresses are required: one for the network adapters on the private network, one for the network adapters on the public network, and one for the virtual server itself. A single static cluster IP address is required for the entire cluster.
- A domain user account for Cluster service (all nodes must be members of the same domain).
- Each node should have at least two network adapters—one for connection to the public network and the other for the node-to-node private cluster network. If only one network adapter is used for both connections, the configuration is unsupported. A separate private network adapter is required for HCL certification.

Shared Disk Requirements

Note: Do not let more than one node access the shared storage devices at the same time until Cluster service is installed on at least one node and that node is online. This can be accomplished through selective storage presentation, SAN switch zoning, or having only one node online at all times.

- All software components listed in the SAN connection tool must be installed and the fiber cables attached to the HBA(s) before the cluster installation is started.
- All shared disks, including the quorum disk, must be accessible from all nodes.
- All shared disks must be configured as basic (not dynamic).
- All partitions on the disks must be formatted as NTFS.

Cluster Installation

During the installation process, nodes will be shut down and rebooted. These steps are necessary to guarantee that the data on disks that are attached to the shared storage bus is not lost or corrupted. This can happen when multiple nodes try to simultaneously write to the same disk that is not yet protected by the cluster software.

Use Table 14 to determine which nodes and storage devices should be presented during each step.

Table 14: Power Sequencing for Cluster Installation

Step	Node 1	Additional Nodes	Storage	Comments
Setting Up Networks	On	On	Not Presented	Verify that all storage devices on the shared bus are not presented. Power on all nodes.
Setting up Shared Disks	On	Off	Presented	Shutdown all nodes. Present the shared storage, then power on the first node.
Verifying Disk Configuration	Off	On	Presented	Shut down first node, power on next node. Repeat this process for all cluster nodes.
Configuring the First Node	On	Off	Presented	Shutdown all nodes; power on the first node.
Configuring additional Nodes	On	On	Presented	Power on the next node after the first node is successfully configured. Complete this process for all cluster nodes.
Post-installation	On	On	Presented	At this point all cluster nodes should be on.

To configure the Cluster service on the HP StorageWorks NAS server, an account must have administrative permissions on each node. All nodes must be member servers within the same domain. It is not acceptable to have a mix of domain controllers and member servers in a cluster.

Setting Up Networks

Each cluster node requires at least two network adapters—one to connect to a public network, and one to connect to a private network consisting of cluster nodes only.

The private network adapter establishes node-to-node communication, cluster status signals, and cluster management. Each node's public network adapter connects the cluster to the public network where clients reside.

Verify that all network connections are correct, with private network adapters connected to other private network adapters only, and public network adapters connected to the public network.

Configure the Private Network Adapter

The following procedures are Best Practices provided by Microsoft and should be configured on the private network adapter.

- On the General tab of the private network adapter, ensure that only TCP/IP is selected
- Ensure that the **Register this connection's address in DNS** is not selected in the DNS tab under advanced settings for the private network adapter
- Ensure that the Link Speed and Duplex is set to 100Mps/Full Duplex under the advanced tab for the Ethernet card used for the private network adapter

Configure the Public Network Adapter

While the public network adapter's IP address can be automatically obtained if a DHCP server is available, this is not recommended for cluster nodes. HP strongly recommends setting static IP addresses for all network adapters in the cluster, both private and public. If IP addresses are obtained via DHCP, access to cluster nodes could become unavailable if the DHCP server goes down. If DHCP must be used for the public network adapter, use long lease periods to assure that the dynamically assigned lease address remains valid even if the DHCP service is temporarily lost. In all cases, set static IP addresses for the private network connector. Keep in mind that Cluster service will recognize only one network interface per subnet.

Rename the Local Area Network Icons

HP recommends changing the names of the network connections for clarity. The naming will help to identify a network and correctly assign its role.

Verifying Connectivity and Name Resolution

To verify name resolution, ping each node from a client using the node's machine name instead of its IP number.

Verifying Domain Membership

All nodes in the cluster must be members of the same domain and able to access a domain controller and a DNS Server.

Setting Up a Cluster User Account

The Cluster service requires a domain user account under which the Cluster service can run. This user account must be created before installing Cluster service, because setup requires a user name and password. This user account should not belong to a user on the domain. This user account will need to be granted administrator privileges.

About the Quorum Disk

When configuring the Quorum disk only one node should be powered on. All other potential cluster nodes must be powered off.

The quorum disk is used to store cluster configuration database checkpoints and log files that help manage the cluster. The quorum disk must be a shared disk resource. HP makes the following quorum disk recommendations:

Note: Use the WebUI Disks tab to configure the quorum disk resource.

- Create a small partition [A minimum of 50 megabytes (MB) to be used as a quorum disk. HP recommends a quorum disk to be 500 MB.]
- Dedicate a separate disk resource for a quorum disk. As the failure of the quorum disk would cause the entire cluster to fail, it is strongly recommended that the disk resource be a RAID 1 configuration.

During the Cluster service installation, a drive letter must be provided for the quorum disk. HP recommends the drive letter Q for the quorum disk.

Configuring Shared Disks

Use the WebUI to configure additional shared disk resources. Verify that all shared disks are formatted as **NTFS** and are designated as **Basic**.

Additional shared disk resources will be automatically added into the cluster as physical disk resources during the installation of cluster services. Each physical disk resource will reside in its own cluster group.

Verifying Disk Access and Functionality

Write a file to each shared disk resource to verify functionality.

At this time, shut down the first node, power on the next node and repeat the Verifying Disk Access and Functionality step above for all cluster nodes. When it has been verified that all nodes can read and write from the disks, turn off the cluster nodes and power on the first, and then continue with this guide.

Install Cluster Service Software

Clustering is installed by default. It is necessary to configure the cluster by launching Cluster Administrator. Follow the steps in the next section to configure the cluster. It is possible to add seven additional cluster nodes for an eight node cluster. Refer to the associated Storage Array documentation to determine the number of cluster nodes that are supported by the specific array in use under Windows Storage Server 2003.

Creating a Cluster

From the WebUI, select the **Cluster** tab:

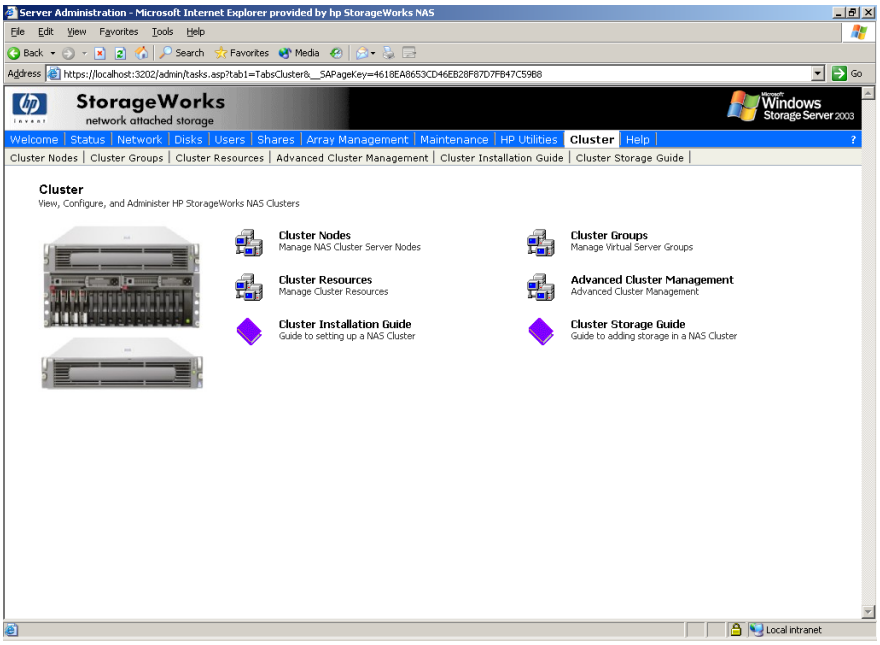

Figure 118: Cluster tab

1. Select **Advanced Cluster Management** to launch a Remote Desktop session.
2. Log into the Remote Desktop session.
3. Click **OK** when the error message regarding a cluster failure is displayed.
4. Select **File > New > Cluster**.
5. In the Welcome to the new server cluster window select **Next**.
6. In the New Server Cluster Wizard window, select the domain in which the cluster will be created and enter the name for the cluster. Select **Next**.
7. In the New Server Cluster Wizard window, enter the computer name to be the first node in the cluster and select **Next**.

The next step runs a pre-configuration analysis. This procedure analyzes and verifies the hardware and software configuration and identifies potential problems. A comprehensive and easy-to-read report will be created listing any potential configuration issues before the cluster is created.

1. Select the details tab to see a list of the items analyzed and any potential issues there may be with the cluster configuration.
2. If there are any issues fix the issues with the suggestion provided in the details tab and then select **Re-analyze**.

   Some issues that may occur are:

   - No shared disk for the Quorum disk. A shared disk must be created with a NTFS partition at least 50 MB in size
   - Use of DHCP addresses for network connections. All Network adapters must be configured with static IP addresses in a cluster configuration
   - Service for Macintosh and Service for NetWare are not supported in a cluster configuration
   - Dynamic Disks are not supported in a cluster configuration
   - Errors will appear on a network adapter that is not configured or does not have an active link. If the network adapter is not going to be used it should be disabled
3. After all issues have been resolved click **Next** to continue.
4. In the New Server Cluster Wizard window, enter the Cluster IP address and click **Next**.
5. In the New Server Cluster Wizard window, enter the login information for the domain account under which the cluster service will be run. Click **Next**.
6. In the New Server Cluster Wizard window, review the proposed cluster configuration and click **Next**.

Note: It is possible to change the Quorum disk by selecting the Quorum button. This will display a list of available disks that can be used for the Quorum disk. Select the appropriate disk and select OK to continue.

---

7. In the New Server Cluster Wizard window, select **Next** to create the cluster.

After configuration is complete the following message is displayed: `You have successfully completed the New Server Cluster Wizard.` Click **Finish** to close the wizard.

Adding Nodes to a Cluster

Note: Only the Quorum disk should be accessible by the new node. The new node should not have access to the LUNs in the cluster until after it has joined the cluster. After the node has joined the cluster, the LUNs may be presented to the new node. Move the physical disk resources over to the new node to confirm functionality.

---

1. Connect to the WebUI of a node that is a member of the cluster. Select the **Cluster** tab, and then select **Cluster Nodes**.

2. Ensure that the additional node has access to only the quorum LUN utilized as the cluster quorum disk.

Caution: Presenting other LUNs to the non-clustered system could lead to data corruption.

3. Click **Add New Node**.
4. Enter the name of the node and specify the password for the cluster service account. Select **OK** to continue.

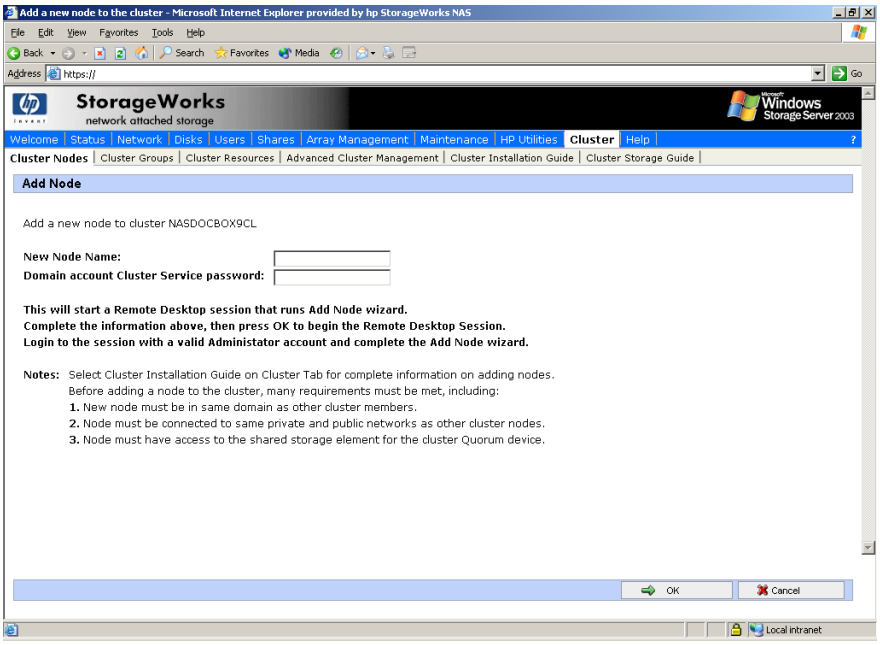

Figure 119: Adding a new node

5. In the Add Nodes Wizard window, click **Next**.
6. Specify the domain and select **Next**.
7. In the Add Nodes Wizard window, confirm the name of the node joining the cluster and select **Next** to continue.
8. The next screen analyzes the configuration to determine the cluster configuration. Potential configuration errors are displayed. Fix any potential errors and select **Re-analyze**. Click **Next** to continue.
9. In the Add Nodes Wizard window, enter the password for the cluster account and click **Next** to continue.
10. The next screen displays a proposed cluster configuration summary. Confirm that all settings are correct and select **Next** to join the cluster.
11. Click **Next** and then **Finish** to complete the cluster wizard.

After the node has successfully joined the cluster present all additional storage LUNS to the node. Please refer to the Cluster Storage Guide located in the Cluster tab for additional information on adding and configuring additional physical disk resources.

Geographically Dispersed Clusters

Cluster nodes can be geographically dispersed to provide an additional layer of fault tolerance. Geographically dispersed clusters are also referred to as stretched clusters.

The following rules must be followed with geographically dispersed clusters:

- A VLAN connection with latency of 500 milliseconds or less ensures that cluster consistency can be maintained. If the VLAN latency is over 500 milliseconds, the cluster consistency cannot be easily maintained.
- All nodes must be on the same subnet.

HP Storage Works NAS Software Updates

After cluster installation has been completed, run the cluster update tool located in the Cluster Install Guide in the WebUI. The Cluster Installation Guide is located under the Cluster tab. This must be done after the cluster installation is complete. It is necessary to run this tool on all cluster nodes.

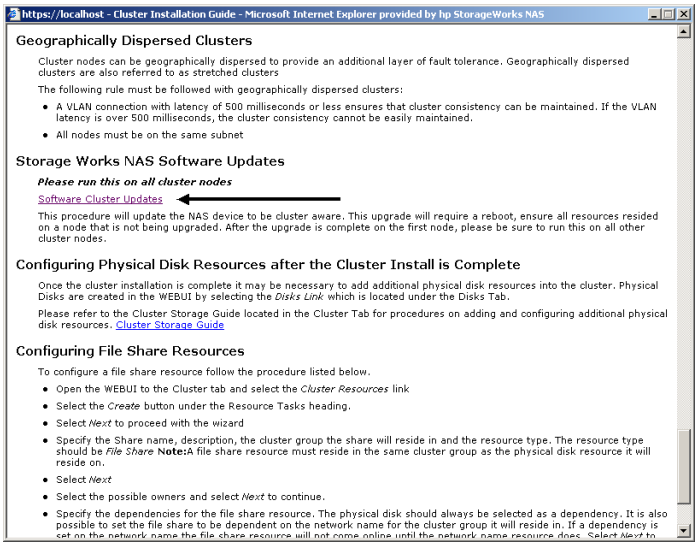

Figure 120: Cluster update tool

This completes the initial cluster installation.

Cluster Groups and Resources, including File Shares

Management tasks for a cluster include creating and managing cluster resources and cluster groups. The Cluster Administrator tool provides complete online help for all cluster administration activities. Cluster resources are created and then assigned to logical, organizational groups. Ownership of these groups should be assigned in a balanced arrangement between the server nodes, distributing the processing load between the two nodes.

Cluster resources include administrative types of resources as well as file shares. The following paragraphs include overview and planning issues for cluster groups, cluster resources, and clustered file shares.

Creating and managing these resources and groups must be managed through Cluster Administrator, available from the Cluster tab of the WebUI. Complete online help for creating the various cluster objects is available in the Cluster Administrator tool.

Cluster Group Overview

A default cluster group is automatically created when the cluster is first created. This default cluster group contains an Internet Protocol (IP) Address resource, a Network Name resource, and the Quorum disk resource. When the new cluster is created, the (IP) address and the cluster name that were specified during setup are set up as the IP address and network name of this default cluster group.

Caution: Do not delete or rename the Cluster Group or IP Address. Doing so will result in losing the cluster and will require reinstallation of the cluster.

When creating groups, the administrator's first priority is to gain an understanding of how to manage the groups and their resources. Administrators may choose to create a resource group and a virtual server (IP Address resource and Network Name resource) for each node that will contain all resources owned by that node, or the administrator may choose to create a resource group and virtual server for each physical disk resource. Additionally, the administrator should try to balance the load of the groups and their resources on the cluster between the two nodes.

Node Based Cluster Groups

Creating only one resource group and one virtual server for each node facilitates group and resource administration. This setup allows administrators to include all file share resources under one group. Clients access all of the resources owned by one node through a virtual server name.

In node based cluster groups, each group has its own network name and IP address. The administrator decides on which node to place each physical disk resource. This configuration provides a very coarse level of granularity. All resources within a group must remain on the same node. Only two IP addresses and network names are required. This configuration creates less overhead for resource and network administration. A possible disadvantage of this approach is that the resource groups can potentially grow large when many file shares are created.

Load Balancing

The creation of separate cluster groups for each virtual server provides more flexibility in balancing the processing load on the cluster between the two nodes. Each cluster group can be assigned to a cluster node with the preferred owner parameter. For example, if there are two cluster groups, the cluster could be set up to have the first cluster group owned by node A and the second cluster group owned by node B. This allows the network load to be handled by both devices simultaneously. If only one cluster group exists, it can only be owned by one node and the other node would not serve any network traffic.

Cluster Resource Overview

Hardware and software components that are managed by the cluster service are called cluster resources.

Resources represent individual system components. These resources are then organized into groups and managed as a group.

Some resources are created automatically by the system and other resources must be set up manually.

Resource Types:

- IP Address resource
- Cluster name resource
- Cluster Quorum disk resource
- Physical Disk resource
- Virtual server name resources
- CIFS file share resources
- NFS file share resources

File Share Resource Planning Issues

CIFS and NFS are cluster aware protocols that support the Active/Active cluster model, allowing resources to be spread out and processed on both nodes at the same time. For example, some NFS file share resources can be assigned to a group owned by a virtual server for NodeA and additional NFS file share resources can be assigned to a group owned by a virtual server for NodeB.

Configuring the file shares as cluster resources provides for high availability of file shares. Because the resources are placed into groups, ownership of the files can easily move from one node to the other, as circumstances require. If the cluster node owning the group of file shares should be shut down or fail, the other node in the cluster will begin sharing the directories until the original owner node is brought back on line. At that time, ownership of the group and its resources can be brought back to the original owner node.

Resource Planning

1. Create at least one virtual server for each node in the cluster.

   A virtual server is a resource group consisting of an IP Address resource and a Network Name resource. Ownership of these virtual servers should be assigned to the different server nodes. In addition to providing load balancing capabilities, the virtual server allows for the transition of group resources in failover situations.

2. Create a virtual server group for each node in the cluster.

   Cluster resource groups are used to balance the processing load on the servers. Distribute ownership of the groups between the virtual servers.

3. For NFS environments, configure the NFS server.

   NFS specific procedures include entering audit and file lock information as well as setting up client groups and user name mappings. These procedures are not unique to a clustered deployment and are detailed in the “Microsoft Services for NFS” chapter. Changes to NFS setup information are automatically replicated to all nodes in a cluster.

4. Create the file share resources.

   In a clustered environment, file shares are created as a type of cluster resource. Creating cluster resources and file shares is documented later in this chapter.

5. Assign ownership of the file share resources to the resource groups.

   a. Divide ownership of the file share resource between the resource groups, which are in turn distributed between the virtual servers, for effective load balancing.

   b. Make sure that the physical disk resource for this file share is also included in this group.

   c. Make sure that the resources are dependent on the virtual servers and physical disk resources from which the file share was created.

Permissions and Access Rights on Share Resources

File Share and NFS Share permissions must be managed via the Cluster Administrator tool versus the individual shares on the file system themselves via Windows Explorer. Administering them through the Cluster Administrator tool allows the permissions to migrate from one node to other. In addition, permissions established using Explorer will be lost once the share is failed or taken offline. To access the permissions, see "Setting Permissions for a SMB File Share" and "Setting Permissions for an NFS Share."

NFS Cluster Specific Issues

In addition to the user name mapping best practices outlined in the "Microsoft Services for NFS" chapter, there are additional recommendations.

For convenience, all suggestions are listed below:

- Back up user and group mappings

  To avoid loss of complex advanced mappings in the case of a system failure, back up the mappings whenever the mappings have been edited or new mappings have been added.

- Map consistently

  Groups that are mapped to each other should contain the same users and the members of the groups should be properly mapped to each other to ensure proper file access.

- Map properly

  — Valid UNIX users should be mapped to valid Windows users.

  — Valid UNIX groups should be mapped to valid Windows groups.

  — Mapped Windows user must have the **Access this computer from the Network privilege** or the mapping will be squashed.

  — The mapped Windows user must have an active password, or the mapping will be squashed.

- In a clustered deployment, create user name mappings using domain user accounts.

  Because the security identifiers of local accounts are recognized only by the local server, other nodes in the cluster will not be able to resolve those accounts during a failover. Do not create mappings using local user and group accounts.

- In a clustered deployment, administer user name mapping on a computer that belongs to a trusted domain.

  If NFS administration tasks are performed on a computer that belongs to a domain that is not trusted by the domain of the cluster, the changes are not properly replicated among the nodes in the cluster.

- In a clustered deployment, if PCNFS password and group files are being used to provide user and group information, these files must be located on each node of the system.

  Example: If the password and group files are located at *c:\maps* on node 1, then they must also be at *c:\maps* on node 2. The contents of the password and group files must be the same on both nodes as well.

  These password and group files on each server node must be updated periodically to maintain consistency and prevent users or groups from being inadvertently squashed.

Non Cluster Aware File Sharing Protocols

Services for Macintosh (SFM), File and Print Services for NetWare, HTTP file sharing protocols are not cluster aware and will experience service interruption if installed on a clustered resource during failover events of the resource. Service interruptions will be similar to those experienced during a server outage. Data that has not been save to disk prior to the outage will experience data loss. In the case of SFM, it is not supported because SFM maintains state information in memory. Specifically, the Macintosh volume index is located in paged pool memory. Using SFM in clustered mode is not supported and may result in data loss similar in nature to a downed server should the resource it is based on fails over to the opposing node.

Creating a New Cluster Group

To create a cluster group:

1. Open the WebUI to the **Cluster** tab and select **Cluster Groups**.

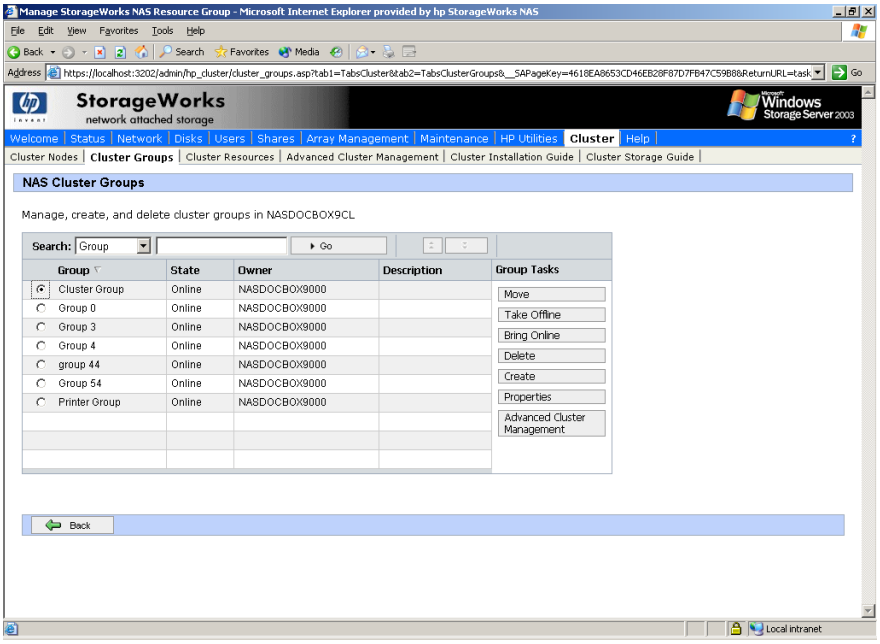

Figure 121: Cluster Groups page

2. Select **Create** to create a new group.
3. Specify the properties for the new cluster group and select **OK** to create the cluster group.

Adding New Storage to a Cluster

Present the new storage to one node in the cluster. This can be accomplished through selective storage presentation or through SAN switch zoning.

Open the WebUI and navigate to the Disks tab. Select the Disks link under the disks tab. Select the disk which needs to be configured from the list of available disks and select Create New Volume. Follow the steps in the wizard to create the new volume. The LUN needs to be configured as a basic disk with a NTFS file system.

Note: If the disk does not appear in the list of available disks on the Manage Disks page then select Rescan to rescan for new disks and refresh the page.

Open the WebUI and select the **Cluster** tab. Follow the procedures listed below to create a physical disk resource.

See the Cluster Storage Guide on the WebUI Cluster tab for detailed information on adding storage elements into the cluster.

Creating Physical Disk Resources

A physical disk resource must reside within a cluster group. An existing cluster group can be used or a new cluster group must be created. See “Creating a Cluster Group” earlier in this chapter.

To create a physical disk resource:

1. In the **Cluster** tab select **Cluster Resources**.

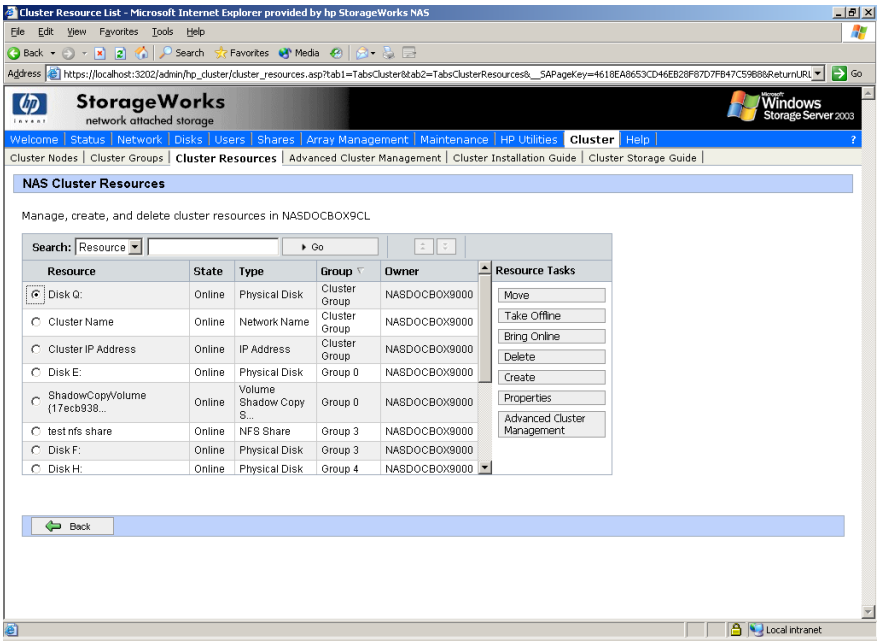

Figure 122: Cluster Resources page

2. Click **Create**.
3. On the Welcome Page select **Next**.

4. Specify a name for the cluster resource and enter a description for the resource.
5. Select the Cluster group the physical disk will reside in.
6. Select Physical Disk as the resource type and select **Next**.
7. Select the Possible Owners and select **Next**.
8. Set the dependencies and select **Next**.

Note: Physical disk resources usually do not have any dependencies set.

---

9. Specify the available disk resource and select **Next**.
10. Review the configuration and select **Finish** to create the physical disk resource.
11. After the resource is created it is necessary to bring it online. In the **Cluster Resources** page, select the resource and select Bring Online
12. Select **OK** on the Bring a Resource Online page to bring the new physical disk resource online.
13. Present the LUN to the additional cluster nodes.
14. Move the physical disk resource to the other nodes to confirm functionality.

To move a resource:

1. Under the **Cluster** tab in the WebUI, select **Cluster Groups**.
2. Select the group and select **Move**.
3. Specify the new location for the group and select **OK**.

Note: In multi node clusters it will be necessary to specify the node to move the group to. When a cluster group is moved to another node all resources in that group will be moved.

Note: When a physical disk resource is owned by a node the disk will appear as a unknown unreadable disk to all other cluster nodes. This is a normal condition. When the physical disk resource moves to another node the disk resource will then become readable.

Creating File Share Resources

To create a file share resource:

1. Open the WebUI to the **Cluster** tab and select **Cluster Resources**.
2. Click **Create**.
3. Select **Next** to proceed with the wizard.
4. Specify the Share name, description, the cluster group the share will reside in and the resource type. The resource type should be File Share.

Note: A file share resource must reside in the same cluster group as the physical disk resource it will reside on.

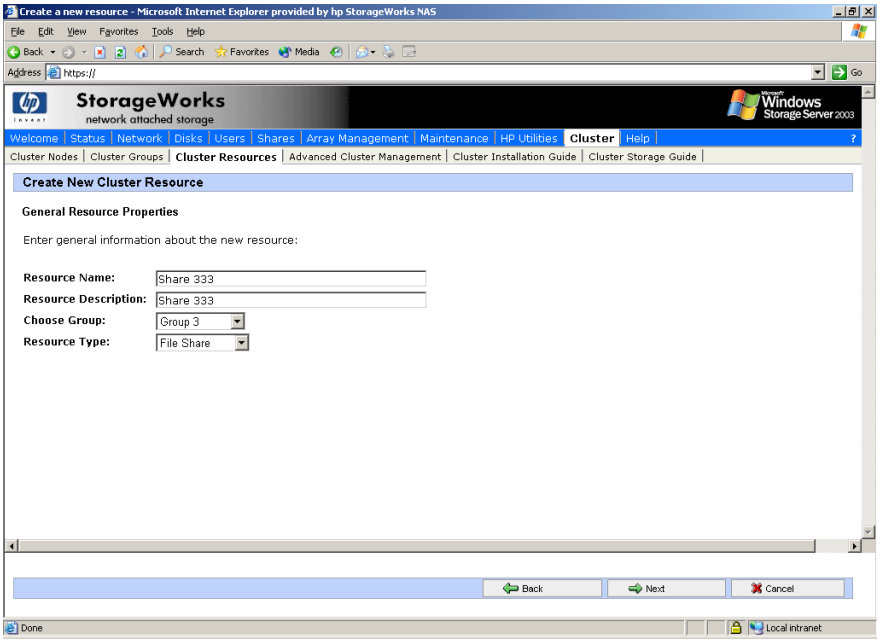

Figure 123: Creating a file share resource

5. Click **Next**.
6. Select the possible owners and click **Next** to continue.
7. Specify the dependencies for the file share resource. The physical disk should always be selected as a dependency. It is also possible to set the file share to be dependent on the network name for the cluster group it will reside in. If a dependency is set on the network name the file share resource will not come online until the network name resource does. Select **Next** to continue.

Note: The physical disk resource specified in this step must reside in the same cluster group as specified in the beginning of this wizard.

---

8. Specify the share name, path, and user limit and click **Next** to continue.

9. Review the configuration and click **Finish** to create the share.
10. After the resource is created it is necessary to bring it online. In the Cluster Resources page, select the resource and select **Bring Online**.
11. Click **OK** on the Bring a Resource Online page to bring the new file share resource online.

Setting Permissions for a SMB File Share

When a share resource is created via the WebUI and brought online, the default permission is set to: Everyone=Read-Only.

To change the default permissions:

1. From the **Cluster** tab, click **Advanced Cluster Management**.
2. Log into Remote Desktop.
3. Click the group.
4. Right-click the resource, then click **Properties**.

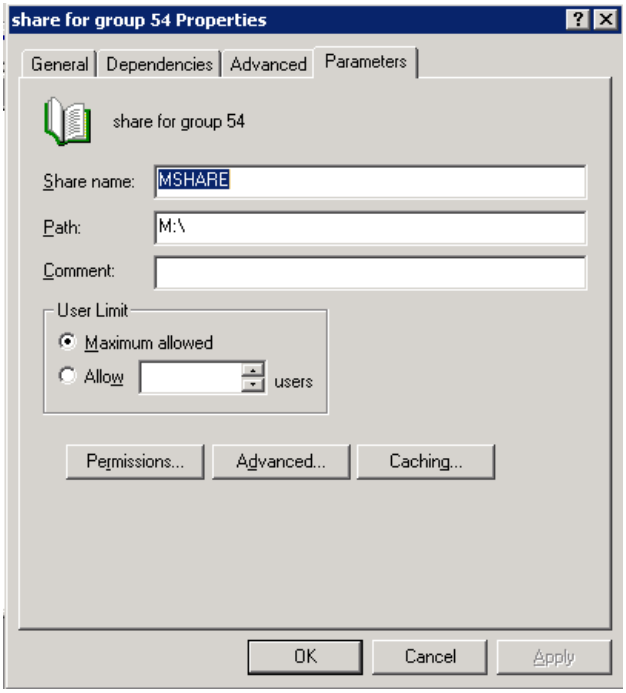

Figure 124: Resource parameters for SMB file share

5. Click the **Parameters** tab.
6. Click **Permissions**.

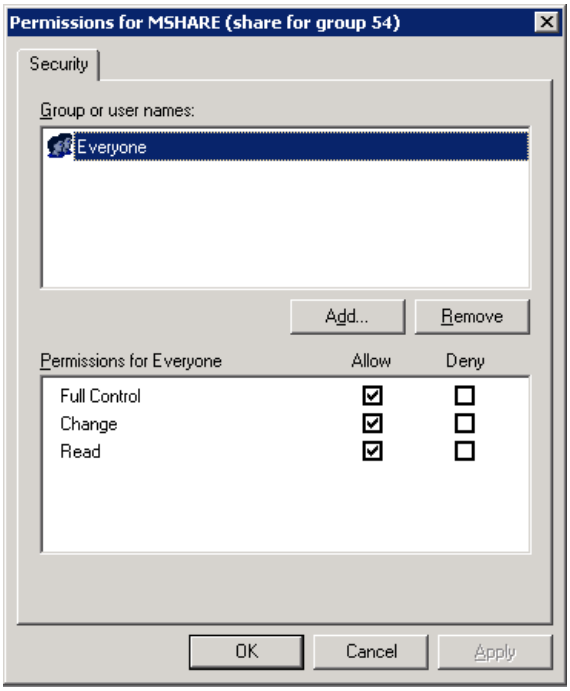

Figure 125: Set resource permissions

7. Set the permissions, then click **OK**.

Creating NFS Share Resources

To create an NFS share resource:

1. Open the WebUI to the **Cluster** tab and select **Cluster Resources**.
2. Click **Create**.
3. Select **Next** to proceed with the wizard.
4. Specify the name, description, the cluster group the share will reside in and the resource type. The resource type should be NFS Share.
5. Click **Next**.
6. Select the possible owners and click **Next** to continue.
7. Specify the dependencies and click **Next** to continue.
8. Specify the share name, path, Share Root Only or Share Sub-directories only, encoding, anonymous access, and anonymous UID/GID and click **Next** to continue.
9. Review the configuration and click **Finish** to create the NFS share.
10. After the resource is created it is necessary to bring it online. In the Cluster Resources page, select the resource and select **Bring Online**.
11. Click **OK** on the Bring a Resource Online page to bring the new resource online.

Setting Permissions for an NFS Share

When a share resource is created via the WebUI and brought online, the default permission is set to: Everyone=Read-Only.

To change the default permissions:

1. From the **Cluster** tab, click **Advanced Cluster Management**.
2. Log into Remote Desktop.
3. Click the group.
4. Right-click the resource, then click **Properties**.

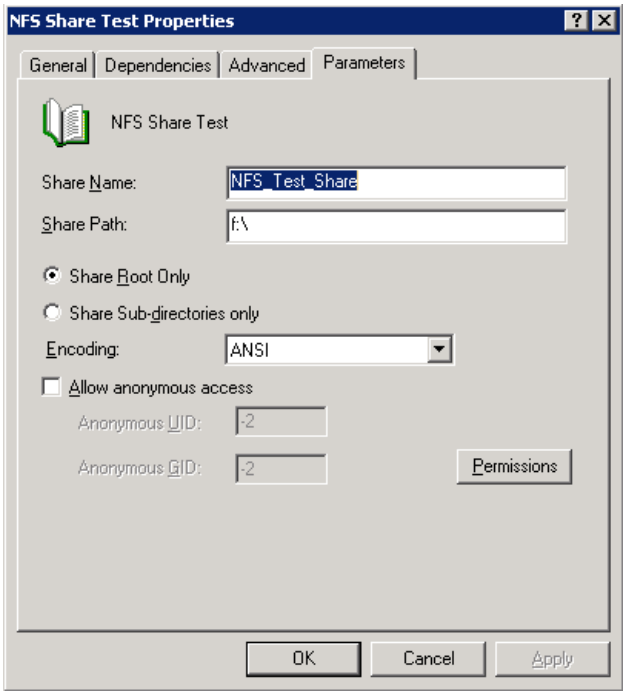

Figure 126: NFS Share Resource parameters

5. Click the **Parameters** tab.
6. Click **Permissions**.

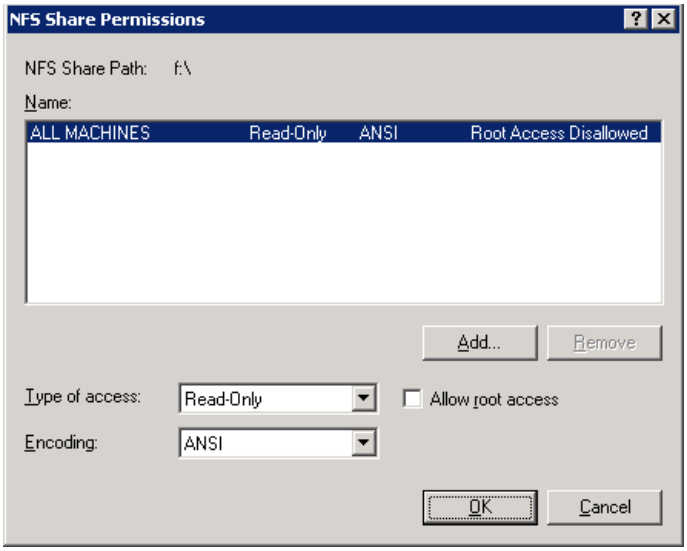

Figure 127: Set NFS Share resource permissions

7. Set the permissions, then click **OK**.

Creating IP Address Resources

1. Open the WebUI to the **Cluster** tab and select **Cluster Resources**.
2. Click **Create**.
3. Select **Next** to proceed with the wizard.
4. Specify the name, description, the cluster group the resource will reside in and the resource type. The resource type should be IP Address.
5. Click **Next**.
6. Select the possible owners and click **Next** to continue.
7. Specify the dependencies and click **Next** to continue.
8. Specify the address, subnet mask, network and whether to enable NetBIOS for this address and click **Next** to continue.

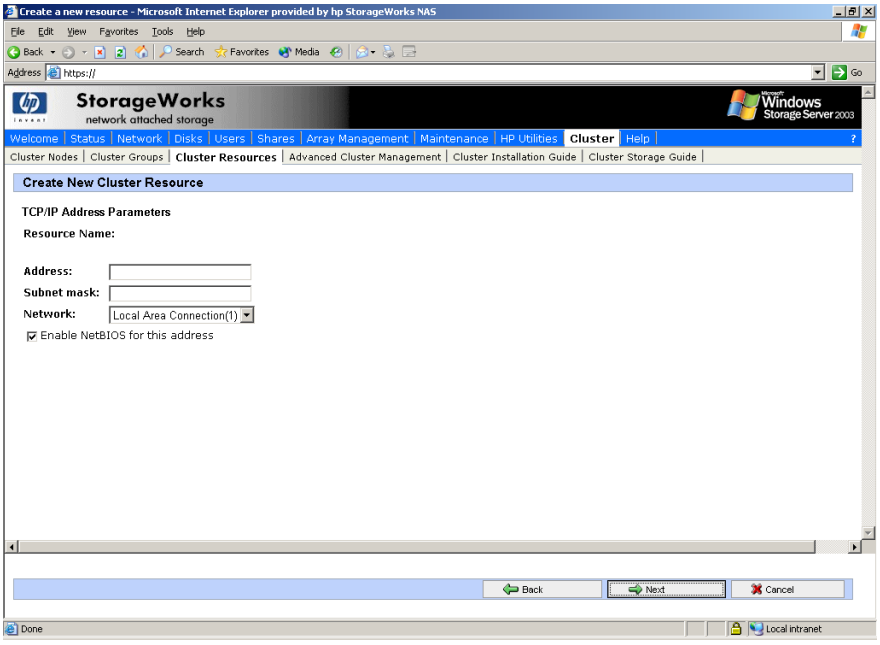

Figure 128: Creating an IP address resource

9. Review the configuration and click **Finish** to create the resource.
10. After the resource is created it is necessary to bring it online. In the Cluster Resources page, select the resource and select **Bring Online**.
11. Click **OK** on the Bring a Resource Online page to bring the new resource online.

Creating Network Name Resources

1. Open the WebUI to the **Cluster** tab and select **Cluster Resources**.
2. Click **Create**.
3. Select **Next** to proceed with the wizard.
4. Specify the name, description, the cluster group the resource will reside in and the resource type. The resource type should be Network Name.
5. Click **Next**.
6. Select the possible owners and click **Next** to continue.
7. Specify the dependencies and click **Next** to continue.

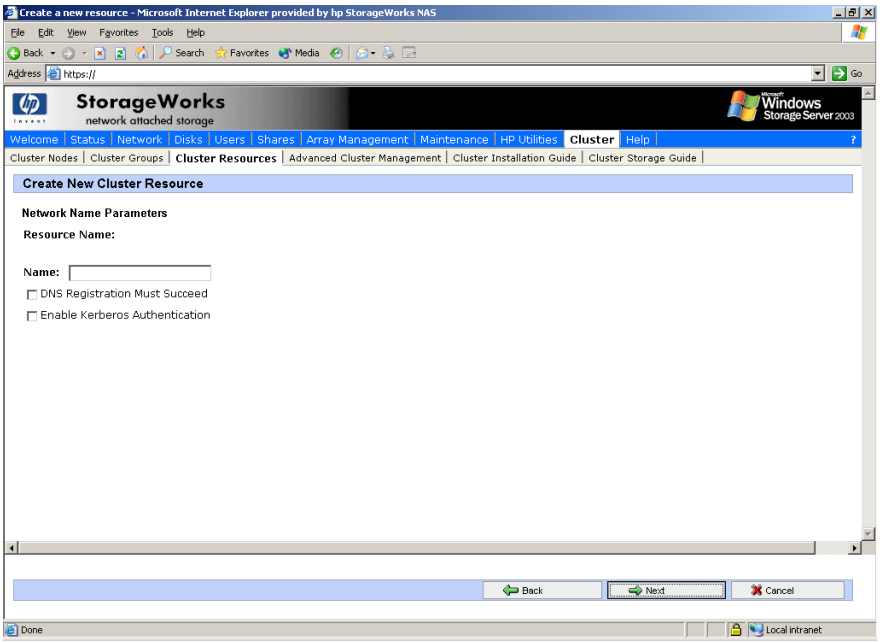

Figure 129: Network Name Parameters

8. Select whether or not DNS registration must succeed and whether to enable kerberos authentication and click **Next** to continue.
9. Review the configuration and click **Finish** to create the resource.
10. After the resource is created it is necessary to bring it online. In the Cluster Resources page, select the resource and select **Bring Online**.
11. Click **OK** on the Bring a Resource Online page to bring the new resource online.

Basic Cluster Administration Procedures

- Failing over and failing back
- Restarting one cluster node
- Shutting down one cluster node
- Powering down all cluster nodes
- Powering up all cluster nodes

Failing Over and Failing Back

As previously mentioned, when a node goes offline, all of the resources dependent on that node are automatically failed over to another node. Processing continues, but in a reduced manner because all operations must be processed on the remaining node(s). In clusters containing more than two nodes additional failover rules can be applied. For instance, groups can be configured to failover different nodes to balance the additional work load imposed by the failed node. Nodes can be excluded from the possible owners list to prevent a resource from coming online on a particular node. Lastly the preferred owners list can be ordered, to provide an ordered list of failover nodes. Using these tools, the failover of resources can be controlled with in a multimode cluster to provide a controlled balanced failover methodology that balances the increased work load.

Because operating environments differ, the administrator must indicate whether the system will automatically fail the resources (organized by resource groups) back to their original node or will leave the resources failed over, waiting for the resources to be moved back manually. See “Managing Cluster Resource Groups” for information on allowing or preventing failback and moving these resources from one node to another.

Note: If the NAS server is not set to automatically fail back the resources to their designated owner, the resources must be moved back manually each time a failover occurs. See “Managing Cluster Resource Groups” for information on overriding this default setting.

Restarting One Cluster Node

Caution: Restarting a cluster node should be done only after confirming that the other node(s) in the cluster are functioning normally. Adequate warning should be given to users connected to resources of the node being restarted.

Attached connections can be viewed through the NAS Management Console on the NAS Desktop using Terminal Services. From the NAS Management Console, select File Sharing, Shared Folders, and Sessions.

The physical process of restarting one of the nodes of a cluster is the same as restarting a NAS device in single node environment. However, additional caution is needed.

Restarting a cluster node causes all file shares served by that node to fail over to the another node(s) in the cluster based on the failover policy in place. Until the failover process completes, any currently executing read and write operations will fail. Other node(s) in the cluster will be placed under a heavier load by the extra work until the restarted node comes up and the resources are moved back.

Shutting Down One Cluster Node

Caution: Shutting down a cluster node must be done only after confirming that the other node(s) in the cluster are functioning normally. Adequate warning should be given to users connected to resources of the node being shutdown.

Shutting down a cluster node causes file shares served by that node to fail over to the other node(s). This will cause any currently executing client read and write operations to fail until the cluster failover process completes. The other node(s) will be placed under a heavier load by the extra work until the second node is powered up and rejoins the cluster.

Powering Down the Cluster

The power down process for the NAS cluster is similar to the process for a single node, but with the cluster, extra care must be taken with the storage subsystem and the sequence of the shutdown.

The power down process is divided into two main steps:

1. Shutting down the cluster nodes
2. Removing power from the cluster nodes

The sequence of these steps is critical. The devices must be shut down before the storage subsystem. Improperly shutting down the nodes and the storage subsystem causes corruption and loss of data.

Caution: Before powering down the cluster nodes, follow the proper shutdown procedure as previously illustrated. See "Shutting Down One Cluster Node." Only one cluster node should be shut down at a time.

Caution: The cluster nodes should never be powered on when the storage subsystem is not available.

Powering Up the Cluster

The power up process for the NAS cluster is more complex than it is for a single node because extra care must be taken with the storage subsystem.

The sequence of the power up steps is critical. Improper power up procedures can cause corruption and loss of data.

Caution: Do not power up the cluster nodes without first powering up the storage subsystem, and verifying it is operating normally.

Nodes should be powered up separately allowing one node to form the cluster prior to powering up the additional node(s). To power up the cluster nodes:

1. After the storage subsystem is confirmed to be operating normally, power up a single node by pressing the power button on the front of the device. Wait for the node to come completely up before powering up the subsequent node.

   If more than one node is powered up at the same time, the first node that completes the sequence will gain ownership of the cluster quorum and will control the cluster database. Designate a particular node as the usual cluster quorum owner by always powering up that node first and letting it completely restart before powering up additional cluster node(s).

2. Power up the additional cluster node(s) by pressing the power button on the front of the device. Each node should be allowed to start fully, prior to starting a subsequent node.

As each node starts, the monitor displays the logon dialog. Background processes will start the cluster service and form the cluster.

Shadow Copies in a Clustered Environment

The creation and management of clustered Shadow Copy resources in a cluster should be performed using the WebUIs by selecting **Disk**, **Shadow Copy** or from the file system by right-clicking on the volume and selecting **Shadow Copy**.

Assuming the underlying disk is part of a cluster, both methods will generate a cluster resource on the cluster that is viewable from Cluster Administrator and the Cluster Resource tab of the WebUI. While the ability to create the Shadow Copy Resource is available in the Cluster Administrator Management Tool, this operation is not supported by Microsoft. The resource may be taken offline/online and managed with the group via all means available.

As recommended in the Shadow Copy chapter, the location of the cache file is recommended on a separate disk from the original data. In this case, a physical disk resource for the cache file disk should be created in the same cluster group as the intended Shadow Copy resource and the volume for which snapshots will be enabled. The resource should be created prior to the establishment of Shadow Copies. The Shadow Copy resource should be dependent on both the original physical disk resource and the physical disk resource that contains the cache file. The update of the Shadow Copy schedule may be done via the Cluster Administrator tool, the WebUI, or the file system.

Creating a Cluster Printer Spooler

Printer spoolers should be created in a separate group dedicated to this purpose for ease of management. For each printer spooler a physical resource is required to instantiate the print spooler resource. In some cases, dedicated physical resources are not available and hence sharing of the physical resource among other members of the group is acceptable, remembering that all members of a group are managed as a unit. Hence, the group will failover and failback as a group.

To create a printer spooler :

1. Create a dedicated group (if desired).
2. Create a physical resource (disk) (if required, see note).
3. Create an IP address resource for the Virtual Server to be created (if required, see note).
4. Create a Virtual Server Resource (Network Name) (if required, see note).

Note: Steps 1-4 may be done via the WebUI interface using the appropriate functions or via the Advanced Cluster Management function and are documented else where in this chapter. If the printer spool resource is added to an existing group with a physical resource, ip address and virtual server resource, steps 1-4 are not required.

---

5. Create a Print Spool resource:
   a. Click the **Cluster** tab.
   b. Select **Advanced Cluster Management**.
   c. Select the group container for the printer spooler.
   d. Right-click and select **Printer Resource**.
   e. Enter the name of the printer resource.
   f. Select all of the appropriate dependent resources (ip address, network name, and physical resource).

g. Select the folder to place the spooler temporary contents and click **Finish**.

h. Close Cluster Administrator.

6. To connect to the Virtual Server Name or IP address created in the steps above:

   a. Select **Start** > **Run** > then type:

      ```
      \\"virtual_server_name or ip address" from the local menu
      ```

   b. A session will open to the virtual server.

7. To add a printer to the virtual server:

   a. Double-click the printers and faxes icon.

   b. Right-click the new screen and select **add printer**. A wizard will start.

   c. Select **create a new port** and click **Next**.

   d. Enter the IP address of the network printer.

   e. Update the Port Name if desired and click **Next**, then **Finish**.

   f. Select the appropriate driver and click **Next**.

   g. If presented with a dialog to replace the driver present, click **keep the driver** and click **Next**.

   h. Name the printer and click **Next**.

   i. Provide a share name for the printer for network access and click **Next**.

   j. Provide location information and comments and click **Next**.

   k. Select **Yes** to print a test page and click **Next**, then click **Finish**.

   l. A dialog will appear regarding the test page. Select the appropriate answer.

The Printer Spool is now a clustered resource.

Remote Access Methods and Monitoring

11

The HP StorageWorks NAS server comes from the factory with full remote manageability. Several methods of remote access are provided:

- Web based user interface
- Remote Desktop
- Integrated Lights-Out Port
  - Features
  - Integrated Lights-Out Port Configuration
  - Using the Integrated Lights-Out Port to Access the NAS server
- Telnet Server
  - Enabling Telnet Server
  - Configuring Telnet Server
- Insight Manager
  - Insight Manager Console
  - Insight Manager Agent Web Interface

These options let administrators use interfaces with which they are already familiar.

Web Based User Interface

The NAS server includes a Web based user interface (WebUI) for the administrator to remotely manage the machine. Of all of the remote access methods, the WebUI is the most intuitive and easiest to learn and use.

The WebUI permits complete system management, including system configuration, user and group management, shares management, Microsoft Services for NFS, and storage management.

To access the WebUI:

1. Launch a Web browser.
2. In the URL field, enter:

   `https://<your NAS b2000 machine name or IP address>:3202/`

Extensive procedural online help is included in the WebUI.

Remote Desktop

The NAS server supports Remote Desktop, with a license for two concurrently running open sessions. Remote Desktop provides the same capabilities as being physically present at the server console.

Use Remote Desktop to access:

- The NAS server desktop
- The NAS Management Console
- A command line interface
- Backup software
- Antivirus programs
- Telnet Server

To access Remote Desktop from the WebUI, select Maintenance, Remote Desktop. For additional procedural information on Remote Desktop, see the “Setup Completion and Basic Administrative Procedures” chapter.

Integrated Lights-Out Port

The following information provides an overview of the Integrated Lights-Out port capabilities. For further information, refer to the *Integrated Lights-Out Port Installation and Users Guide* on the Documentation CD.

The Integrated Lights-Out port is an ASIC-based Web interface that provides remote management for the server.

Regardless of the state of the host operating system or the host CPU, complete capability for the server is available. The Integrated Lights-Out port is independent of the host server and its operating system. The Integrated Lights-Out port provides remote access, sends alerts, and performs other management functions, even when the host server operating system is not responding.

Features

The Integrated Lights-Out port provides the following features:

Note: The remote client console must have a direct browser connection to the Integrated Lights-Out port without passing through a proxy server or firewall.

- Hardware based graphical remote console access
- Remote restart
- Server failure alerting
- Integration with Insight Manager
- Local Area Network (LAN) access through onboard NIC
- Browser support for Internet Explorer 5.50 or later
- Reset and failure sequence replay
- Auto configuration of IP address through domain name system (DNS) or Dynamic Host Configuration Protocol (DHCP)
- Virtual power button

Security Features

- SSL encryption for login and network traffic
- User administration allows capability to define user profiles
- Event generation for invalid login attempts
- Logging of user action in the Event Log

Manage Users Feature

The Manage Users feature allows those with supervisory access to add and delete users or to modify an existing user's configuration. Manage Users also lets the administrator modify:

- User name
- Logon name
- Password
- Simple network management protocol (SNMP) trap IP address
- Receive host OS generated SNMP traps
- Supervisor access
- Logon access
- Remote console access
- Remote server reset access

Manage Alerts Feature

The Manage Alerts feature allows the user to:

- Select alert types received
- Generate a global test alert
- Generate an individual test alert
- Clear pending alerts
- Enable alerts

Refer to the *Integrated Lights-Out Port User Guide* for more information about the Integrated Lights-Out port features and functionality.

Integrated Lights-Out Port Configuration

The Integrated Lights-Out port on the NAS server is initially configured through the Rapid Startup Utility. SNMP is enabled and the Insight Management Agents are preinstalled.

The Integrated Lights-Out port comes with factory default settings, which the administrator can change. Administrators may want to add users, change SNMP trap destinations, or change networking settings. Refer to the *Integrated Lights-Out Port User Guide* for information about changing these settings.

There are several methods for performing Integrated Lights-Out port configuration changes:

- Web interface
- Integrated Lights-Out port configuration utility accessed by pressing **F8** during a system restart

---

Note: You must connect locally with a monitor, keyboard, and mouse to utilize the F8 feature.

- Integrated Lights-Out port access using the default DNS name

Using the Integrated Lights-Out Port to Access the NAS Server

Using the Web interface of a client machine is the recommended procedure for remotely accessing the server:

1. In the URL field of the Web browser, enter the IP address of the Integrated Lights-Out port.

Note: The iLO port can also be accessed from the HP Utilities tab of the WebUI by clicking the remote management link.

---

2. At the Integrated Lights-Out Account Login screen, supply the username and password for the iLO and click **Login**.
3. Click the Remote Console tab. The Remote Console Information screen is displayed.
4. Click the Remote Console choice in the menu on the left side of the screen.
5. Press **Ctrl-Alt-Del** to log in to the console.
6. Supply an administrator username and password. The NAS server desktop is displayed.

Note: The remote desktop feature of the iLO port requires a license key. The key is included with the product inside the Country Kit. See the iLO Advanced License Pack for activation instructions.

Telnet Server

Telnet Server is a utility that lets users connect to machines, log on, and obtain a command prompt remotely. Telnet Server is preinstalled on the NAS server, but must be activated before use.

Caution: For security reasons, the Telnet Server service must be restarted each time the server is restarted.

Enabling Telnet Server

Telnet Server can be enabled in two ways. The first is to use Remote Desktop to access a command line interface and enter the following command:

```
net start tlntsvr
```

The Telnet Server service needs to be enabled prior to running this command. The service can be enabled by opening the services MMC:

1. Select Start, Run, then type services.msc.
2. Locate the Telnet service, right-click on it, then select **Properties**.
3. In the startup type drop-down box, choose **Manual**, and click **OK**.

The second is to open the WebUI:

1. Click **Network**.
2. Click **Telnet**.
3. Check the **Enable Telnet access to this appliance** box.
4. Click **OK**.

Sessions Information

The sessions screen provides the ability to view or terminate active sessions.

HP Insight Manager Version 7

The NAS server is equipped with the latest Insight Management Agents for Servers, allowing easy manageability of the server through HP System Management, HP OpenView, and Tivoli NetView.

Insight Manager is a comprehensive management tool that monitors and controls the operation of HP servers and clients. HP Insight Manager Version 7.0 or later is needed to successfully manage the NAS server using the following components:

- Windows-based console application available on the Insight Manager 7 CD-ROM loaded on a separate client for NAS devices
- Server or client based management data collection agents

Management agents monitor over 1,000 management parameters. Key subsystems make health, configuration, and performance data available to the agent software. The agents act upon that data by initiating alarms in the event of faults. The agents also provide updated management information, such as network interface or storage subsystem performance statistics.

Index

D

E

F

G

H

I

L

M

N

O

P

Q

R

S

T

U

V

W